PLACE: Dobbs Building, Raleigh, North Carolina

DATE: Monday, January 21, 2020

DOCKET NO.: EMP-100, Sub 164

TIME IN SESSION: 1:00 p.m. to 3:12 p.m.

BEFORE: Chair Charlotte A. Mitchell, Presiding

Commissioner ToNola D. Brown-Bland

Commissioner Lyons Gray

Commissioner Daniel G. Clodfelter

Commissioner Kimberly W. Duffley

Commissioner Jeffrey A. Hughes

Commissioner Floyd McKissick, Jr.

IN THE MATTER OF:

Investigation of Energy Storage in North Carolina

Presentation by:

Dr. Jeffrey Taft, Chief Architect for Electric Grid Transmission, Pacific Northwest National Laboratory

and

Dr. Andrew Mills, Research Scientist, Electric Markets and Policy Group, Lawrence Berkeley National Laboratory

Volume 4

P R O C E E D I N G S

CHAIR MITCHELL: Good afternoon, and welcome. I'm Charlotte Mitchell, the Chair of the Utilities Commission, and with me this afternoon are Commissioners ToNola D. Brown-Bland, Lyons Gray, Daniel G. Clodfelter, Kimberly Duffley, Jeffrey Hughes, and Floyd McKissick.

This is the fourth in a series of presentations pursuant to the Commission's September 4th, 2019 Order Initiating Investigation in Docket Number E-100, Sub 164, in which the Commission has initiated a series of educational presentations by invited experts on energy storage related topics.

We're happy to have with us today Dr. Jeffrey Taft and Dr. Andrew Mills. Dr. Taft is the Chief Architect for Electric Grid Transformation at PNNL in Washington state, and Dr. Mills is a Research Scientist in the Electricity Markets and Policy Group at Lawrence Berkeley in California.

Our speakers will be working from slide decks that will be displayed on the monitors here in the hearing room this afternoon. These slides have also been posted on the Commission's website in this docket which is E-100, Sub 164.

Our court reporter is creating a transcript that will be filed in the docket and available on the Commission’s website.

These sessions are structured for the benefit of the Commission’s learning and understanding, and the speakers will be asked to share their expertise and answer the Commission’s questions. People in the audience will not have the opportunity to ask questions, however, if you want to file information in this docket in response to what you hear or if you’d like to suggest other speakers who could appear before the Commission, please file these comments or suggestions in the docket for our future planning purposes.

Gentlemen, if it’s okay, we’d like to ask you all questions as we go, if that’s acceptable.

DR. MILLS: That would be great.

CHAIR MITCHELL: Okay. Thank you. We will do that. Again, we appreciate your preparing this material and spending your time with us today, and look forward to hearing form you.

So I will turn it over to you all. I assume you’ve arranged an order of presentation. Okay.

DR. TAFT: Well, those are my slides, so I guess I’ll start.

CHAIR MITCHELL: Okay. You may start. Thank you very much.

DR. TAFT: I'm Jeffrey Taft from PNNL, and I actually live and work in Pennsylvania. The lab is in the state of Washington, but I live near Pittsburgh. Sometimes I call it Pacific Northeast National Laboratory. And it's in the county -- it's in Washington County in Pennsylvania, so the lab is in the state of Washington. Of course, DOE being our primary sponsor is in Washington, D.C., and I live in Washington County, so when my boss says where are you going to be, I used to say Washington, pick one, you know, whatever you want.

So I'm going to talk about storage in a way that may be a little bit unfamiliar to the Commission today. It's relatively new thinking that came out of work that we've been doing on grid architecture for the last several years under the sponsorship of the U.S. Department of Energy. And to do that, I'm going to start off talking a little bit about that discipline of grid architecture so you can see how we get to some of the answers that we get to and why we think the way we do.

Along the way, some of the ways that we use to reason about that may be useful to the Commission as well; not just the results that we get, but the way that we get there. And I say that because we've worked with more than a dozen state commissions on this type of work, and we know from filings and so on that at least 26 state commissions make use of our work in one way or another, so I think they find it useful. And so we're going to talk a little bit about that, then we'll talk specifically about this idea of storage as being something that you would treat as core infrastructure to the grid as opposed to ancillary services devices. So next slide, please.

One of the biggest problems that we have in dealing with the grid is managing the complexity of it. You know, I want you to appreciate that what you are working on has a level of complexity that's so large we actually have a special name for it in the grid architecture world and the system architecture world. It's called ultra large-scale complexity.

If you look at that -- that illustration there, I made a little bit of a graph. There's a course that's taught on system architecture at MIT, and in that course they talk about complexity of systems, and they give us an example of what they call medium complexity, a refrigerator. So I thought about that and I thought, well, if that's the case, we'll say a kitchen timer is low complexity, and refrigerator is medium complexity, a space shuttle is pretty high complexity. Most people would agree with that. Then way off to the right so far it shouldn't even be on the page is our power grids.

That's the problem with some of this stuff, is that you have this amazing amount of complexity that we've inherited, it's legacy, and we're trying to make changes and understand the nature of the changes and the implications of it. The reason that we have all this complexity is because the grid is made up of a number of different things that are interconnected in complex ways. So next slide, please.

To deal with complexity, we use a discipline that's used in a number of different fields. It's used in electronics, it's used in aerospace, it's used in defense and a lot of places where they deal with complicated systems. It's called system architecture.

Now, architecture itself is a word you probably hear a lot, and to some extent it's overused, sometimes it's misused. But when we talk about it, we're talking about a depiction of a complex system, this kind of abstract. And it gives us the ability to reason about that system without going down into all the details. That's part of the way that we manage the complexity.

So an architecture really has three kinds of things to it. It has what we call black box components, it has structure, and it has externally visible characteristics. Now, when I say black box components, that’s kind of important for the way that we’re going to talk about storage because we do not at this level concern ourself with the internal details of how those things work.

So when we talk about storage at the architectural level, it doesn’t matter whether it’s lithium ion or sodium sulfur or hauling a railway car full of boxes -- rocks up a hill. We don’t care how it works. What we care about is what it looks like from the outside. How much energy does it store, how fast does it go in and out, that kind of thing. So that’s what we mean by black boxes.

Structure is the way things are connected together. So if you’ll think about a block diagram for a second, the boxes are the components, and we’re not going to look inside the boxes. The lines that connect them are the structure, and we focus a lot on that for some very good reasons.

So what we did was we took the discipline of system architecture, and when I came to the lab six years ago we went to the Department of Energy and said, look, we can use this for architecture of our power systems of our grid, so we call that discipline that applies system architecture to the grid, grid architecture. So when you hear me use that phrase, that's what that means.

So it's the application, and one of the things that's really useful about it is it helps us reason about the properties, behavior, implications of change to our grid without having to go down into details and without having to spend a lot of money to find out what's going to happen.

One of the problems that you have with complex systems is that when you make a change somewhere, it's kind of like dealing with a tapestry. If you tug on the thread someplace in the tapestry, it's going to bunch up somewhere else, but you might not know where that is until it happens. Well, with the grid we'd rather know about those things before they happen, so grid architecture is a discipline that helps you understand that stuff.

So it has a lot of different purposes, and some of them are listed there, but the one that I marked in red that's maybe the most important is it helps you manage complexity because complexity is the big hidden bear in the room for understanding all of this stuff, so this discipline helps us think about that. It gives us a lot of other tools as well. Next slide, please.

In the grid architecture work we focus on structure a lot, and this is really important because the grid is composed of a number of different structures. The one that everybody would automatically think about is the electric infrastructure, the circuits, the substations and so on, but there are a lot of other structures we have to deal with as well. One of them is industry structure, and that means the collection of entities, the different kinds of businesses, the different kinds of organizations involved and how they relate to each other and, of course, that's different in different parts of the country. In an area where you have vertical integration, you have a different kind of industry structure than you might have at places where they are restructured and have things like system operators and so on. So all of those different structures, and you can see several different classes of them, they are in the gold boxes, comprise the grid and are interconnected with each other in complicated ways, and that's where all this complexity comes from.

So why do we focus on structure so much? You're going to see that when I talk about storage here in a little bit. Well, in the box there you see two reasons. If you get the structure right, the downstream decisions become a lot simpler to make. Things become much clearer. If you don't get the structure right, you run into a very high risk of stranded assets, stranded investments, unrealized benefits, and we've seen this time and again. You can greatly simplify the problem by thinking about the structure first.

Now, we have inherited a massive amount of structure in our power systems from the 20th century. They were designed a particular way for reasons that were well and good at that time, but the problem, as you know probably as well as anybody, is that we are changing the rules. We are changing the way we want things to work. Some of those changes are, in fact, structural changes. Some of them are impeded by the legacy structure that we've had in the past.

When I talk about storage in a few minutes, I'm going to show you one of the biggest problems with structure in the grid today that storage can address but does not presently address. Next slide, please.

Some of the work we do is very complex, and we use mathematical methods and all that, and it's not my suggestion that we try to turn everybody into architects, but you don’t have to be an architect to use the results of this work. In fact, some of the biggest uptake we get in the use of grid architecture is among regulators. We’ve worked with regulators in a large number of states. We know that our work is used in even more states. And one of the things that we get frequently as feedback from that work is that we help make issues crystal clear. That comes about not because we’re smarter than anybody else; it comes about because of the methods that we use, and we find those methods to be helpful. So next slide, please.

So just to be a little bit clearer about this, we start off with definitions because you will run into an amazing number of terms that people use and throw around without being entirely clear about them. Some of those terms have multiple definitions. Some of them are ambiguous. And we always start off with let’s be clear about what we mean about these various terms. I’m going to show you in a little bit how bad that gets.

And as I mentioned, we focus on structure. We use some foundational principles to deal with all that. And we are driven by things like user requirements, emerging trends, and public policy. We don’t try to determine public policy. In fact, we’re not allowed to from the lab, but we have to think about that.

There are things, though, that we are agnostic to, and this is something that I think can be very helpful for your work. We're specifically agnostic to products and services, so you will not see us saying, well, we should use so-and-so product as part of this architecture for the grid. We are agnostic to business models, so we don't spend time thinking about who makes what money and how they make it. And then we try to be agnostic to a hype cycle, so when something new comes along, there's lots of attention paid to it, you know. Blockchain, might have heard a little bit about that in the last few years. We try not to get caught up in that, and we try to see what things really are going to turn out to be.

And so we use these principles, and there is more that I haven't shown here, to develop what we call reference architectures. They're model architectures, and we've been doing that for DOE for some time now, and those model architectures are available to the public. They're on our grid architecture website, so you can have a look at them, your staff can have a look at them, but they are intended to be instructive. We don't view them as being prescriptive as in, well, you know, here we are from the government, we're going to tell you how to build your grid. It's not like that. They are used to illustrate the concepts that people can adapt for their own purposes. So that's how we think about it, and we're really trying to help manage the complexity and produce the insight that enables anybody, any of those stakeholders, whether it's the regulators, whether it's the product developers, it's the operators, to develop the insight to make great decisions because they're the ones that are best positioned to do that, but we can help them sometimes. Next slide, please.

And in that regard there's this little thing that we refer to as a virtuous circle. In the upper left-hand corner you see objectives there. This -- the setting of objectives in the beginning is just incredibly crucial, and it's amazing to me how many times I've seen people try to jump into grid modernization without being clear about what it is they're trying to achieve.

So we did a bunch of work with the Ohio Commission a while back, and you may have seen their PowerForward work there. I worked with the Commission there. And we actually helped them set up a little process to go through to figure out what they wanted to have for their objectives and how that would flow through to the eventual product that they develop which was their document.

So when you do this in an architectural sense, you start off with these objectives, and they come from things like what are the user’s needs, what are the public policies saying, what are the emerging trends to be dealt with. An emerging trend might be penetration of solar into your power system.

Those objectives are going to imply that you need a certain set of capabilities, and you can compare that to the ones you actually have in your systems now and figure out whether there are any gaps. Once you have understood that, that says you have to have certain kinds of functions, and that implies architectural elements and the properties that come with them, those qualities that result from a system built that way should come back around and support those objectives. If you go around that circle and that loop doesn’t close, something has not been done right and you need to revisit it.

So this is a sort of simple thing that you can do early on in the process when you’re thinking about how you want to give guidance to the utilities, for example, how you want to think about your ratemaking. The utilities can use this when they want to think about their modernization plans previous to actually doing their designs. This model works very well. And the real key is to get the objectives right in the beginning. We find so often that people sort of jump over that or presume that everybody agrees about the objectives when, in fact, maybe they haven't really thought that through. All right. Next slide.

So what I want to do now with that is a preamble and understanding that we have a pretty large discipline and connected body of knowledge around all of that, is talk about some things that we have thought about for how to use storage not as a grid services device, but as embedded into the core infrastructure of the grid. So next slide, please.

So when we do this work, we think about emerging trends and the resulting systemic or crosscutting issues that come from all of that. So some of the things that you've been dealing with that are being dealt with in a lot of parts of the country is the fact that generation which used to be, you know, bulk power system connected and more or less centralized, is now being split into a combination of that plus distributed generation, which is connected at the distribution level and which provides a very different set of challenges for how the system operates. As a result, some of those sources, especially the renewables, are very volatile, and that means we can't really predict or dispatch those, so they behave in a different way, and that creates problems in operating the grid in a balanced and sensible way.

We have in a lot of areas an increasing interdependence and integration with natural gas systems because a lot of generation is being powered by natural gas, and so we've seen situations where that interdependency can be a weakness, but we also see opportunities there related to storage, in particular, to make those things work better.

And then this whole business of ubiquitous communication is an interesting systemic issue, too, and what I mean by that, of course, is the digital communication and connection to the internet and the resultant flexibility and capabilities, but also the resultant vulnerabilities that come about from it as well.

So when we think about this set of issues, and there's a much larger set that we actually work with in our reference architecture work, and that's all on the architectural website, but we think about these issues, we think about two characteristics in particular for the grid; resilience and operational flexibility. So we're going to talk about that a little bit because that will show you why we think about storage the way we do. So next slide, please.

This horrible list was actually compiled by a researcher at Caltech named John Doyle, and it's known as the "ilities list" because a lot of the words end in "ility," like flexibility and reliability and so on. Now, not all of them do, but there are just like 80 terms there, and there are even a few more that have come along since then. And what happens in a lot of cases, and we saw this even with DOE going back five years ago, people would show up and say, well, the grid needs to be flexible and adaptable and adjustable and reliable, and they would throw all that stuff out there and say this is what the grid has to be, as if that is -- represents the objective they're trying to achieve. But unfortunately, most of those terms don't have good definitions, they're not quantifiable, and turning them into something meaningful so that you can say this is what I actually have to do or this is how we have to think about it has proven to be very difficult.

So we deal with that a lot, and most of these, by the way, as far as I know there isn't anybody who tries to deal with all of these at one time. Everybody picks their favorites out of this list, you know, and I could show you a lot of examples of that. Even DOE did this in the beginning. They picked their favorites out of the list and they made a nice big slide, and it was wonderful. They don't use that slide anymore because they figured out it wasn't really helping anybody.

Okay. So what do you -- how do you deal with all of that stuff when you know, you know, instinctively you know there's something you're trying to achieve about making electric power service better and then apply something about the grid and grid modernization, and that all implies something about the use of storage, and then you have these issues that arise that are very specific that you have to deal with, so you've got to sort all that stuff out.

Well, in this slide here, one of the things that we show people is you can deal with a lot of this by recognizing -- first of all, you can make pretty clear definitions of these terms, and we do that, and we post that on our website, and then the resource slides at the end of this deck there are some definitions for some of this. But the more important thing to realize here is that these don't stand in isolation. These are related to each other, and there's a structure into which they fit, and that diagram gives you an illustration of that.

Now, I mentioned that we were going to focus on flexibility and resilience in our discussion about storage here, but you can see when you do that, what happens automatically because of where flexibility is positioned, that you're going to also have an impact on reliability as a result. So there's a lot of confusion about resilience versus reliability. We've done a fair amount of work to help untangle all of that. And I wasn't going to go into that in great depth today, but if you ask questions about it, I will stop and talk about that. What I want to show you is that if you think about flexibility and resilience in terms of the structure of the grid, you come up with some different ways to think about storage and how to apply it. So next slide.

A quick definition, we classify storage into two types, what we call reflexive and transitive. The one that we're concerned about is the reflexive. That means electricity. Electric energy comes from the grid, goes into storage, resides there for a while, goes back into the grid in the form of electricity. Now, there are other forms of storage in which you may take that energy and use it for something else. Maybe you use it for preheating a building, or there's a lot of different ways to use storage. They're all valuable, but we're not going to talk about those other ways. We're going to focus on the one in which electricity goes into storage and goes back in the grid in the form of electricity. We call that reflexive storage. Next slide, please.

And that really -- as a component or an element, there's kind of three pieces to it. There's the core technology that stores the energy, there's a controls and advanced controls mechanism, and there's some kind of fast and flexible interface. Now, because of the way most of these things work, that interface is usually in the form of power electronics, referred to as inverters, so that's why you'll come up with that discussion a lot.

So those three things together represent the kind of storage we're going to talk about, and we're going to talk about it in terms of its key characteristics. So remember back in the beginning of my discussion I said we treat these things as black boxes. So we're not going to talk about battery chemistry or electromechanical devices; we're going to talk about what do you see from the outside? So if you go to the next slide.

There's only a handful of characteristics that you need at the architecture level to think about this, and they're listed there, and they are things like how much energy does it store, and how fast does it go in and out, how much do you lose on the round trip. Those are the kind of things that you would see from the outside. And it doesn't matter how the box works. Those things, if you focus on those, help you think about storage and what it's for and what it's going to do and how you want to use it without getting tangled up in all the details. And that kind of abstraction, if you will, is one of the keys to help dealing with all this complexity. So you don't have to get into all those gory little details that people just love to talk about so much. Next slide, please.

So there's been a lot of thought about how to use storage, and people have come up with a lot of different approaches to it, mostly in terms of grid services. So we've seen a lot of models that say, well, we can have all these different things that it can do. And I was going to bring you two pictures, but because of digital rights management, I didn't bring them.

One of the pictures is, if you've ever seen it, there's a company called Wenger that makes what we call Swiss Army knives. Several years ago they made one just as almost a marketing gimmick and it has 47 blades in it. It's about this wide (indicating), and it has a hundred and some functions and all these -- all these things on it. You couldn't possibly use it, but it has everything they ever made all in one thing. And sometimes that's the way people talk about storage. It's got all these different functions, so we can do all these things with it and it must be really great because of that.

The other picture I was going to show you is something very much simpler. It's a shock absorber. And that's how we're going to talk about storage today, is as if it's a shock absorber.

The funny thing about the way we think about the grid is that we haven't really considered what is the core deficiency in the grid that we need storage for. So people have lots of different applications and lots of different ideas, but getting down to the real essence of why does it matter, what does it do for us in the grid that we don't already have the ability to do is what we try to get at with the grid architecture work.

Now, by the way, a lot of people want to supply a lot of different services with storage, and you've probably seen all of that. We actually made a catalog of grid services. It's on our grid architecture website. If you ever get interested in that, we went through and found all the ones that we knew of, including all the ones that are defined by FERC and everybody else and all the ones that we knew of that were being proposed, and we categorized those in a way so that people could look at them and make some sense out of all of that. There are about 40 of them on there. And to do that we looked at some of the lists that came from places like some of the national labs in Southern California and so on and people were thinking about all this kind of stuff. Okay. Next slide.

So when you think about using storage for ancillary services, you're working at the margins. You're sort of working at the edge. A lot of that stuff is not the core of the grid. And that's one of the things that sort of puzzled us about storage, is why it wasn't being used in a more transformative or ubiquitous way, and part of that, I think, is historical.

If you look at what happened with storage in California, and I realize the structure there is quite different, but they decided they were going to fit storage into the same category as generators, and they said, well, it's just like a generator except sometimes it has a negative output. And the reason they know -- and I know this because I know people that were working on this. The reason they did that was because it was easy to fit into their software and their procedures. But that sort of made a very narrow view of what storage could be, and to some extent that view has proliferated to the point where people sort of take that as the model for storage. Well, it's like a generator except it could have negative output, so we think about it that way. That's sort of missing a lot of the point, unfortunately.

And so while people have used storage to kind of improve reliability of the grid on the margins, they haven't really thought about capitalizing on its main capability, as I said, to think about it as being a shock absorber. Next slide, please.

There's something about the grid that's unique compared to other kinds of complex systems, and I'll tell you that I've been doing -- I'm an electrical engineer by education and experience, but I've been doing architectural work for quite a long time, so looking at a lot of different kinds of systems in different fields, and one thing that struck me about the grid is in almost every other kind of complex system we have some form of buffer. So in communication systems we have bitstreams in which the bits are coming in in irregular bursts, but we maybe want to play them back for display in a regular basis, and so we have a thing called a jitter buffer and it evens out that flow.

In logistic systems we have buffers. They're just called warehouses. They even out the flow between the incoming stuff and the outgoing stuff. In gas and water systems we have them; they're called tanks. Almost everywhere that you look at complex systems you'll see buffers except in the grid. The 20th century grid didn't have so much need for it, also didn't have good ways to do it, but that has changed. And so, you know, we -- the grid doesn't have this inherent springiness or sponginess that other complex systems have built into them, and that -- fundamentally what that does is it decouples these mismatched volatilities.

Well, in the past when generation was dispatchable and load was, frankly, fairly predictable, we didn't have too much of a problem with that. So along comes wind and solar and along comes distributed generation, and we kind of upset that whole model that said that we can have -- you know, we can dispatch generation to be load following and because the load terms are pretty well behaved, that's all going to work real well, and we didn't need storage because we could do that balance, and that's what the balancing authorities do, right? They maintain that balance very finely and they have very clever mechanisms by which they do that.

Well, that's fine until you start to have stochastic sources of generation and you start to have all this unpredictable variability, and that's where things become difficult. So if we think about that, the missing sponginess, the missing shock absorbers interior to the grid are the thing that are really holding us up from being able to think about the grid both in terms of overall resilience and in terms of the flexibility to deal with these changes that are coming about almost organically; proliferation of wind and solar, for example. Next slide, please.

So when I think about storage as a shock absorber for the grid, and there are a variety of things that you would do when you do that; there's a long list here and I don't want to read them to you, but if you look at these, they all have one thing in common, the buffering of variable energy flows is the issue, and storage is actually the answer to dealing with that issue. So that's why I would have shown you a picture of a shock absorber and said this is a model for storage, not the giant Swiss Army knife. So let's talk about how to do that. Next slide.

If we want to take storage and embed it in the grid as opposed to attaching it at the edges so that we have this interior springiness or sponginess to deal with all this variability and to provide us with operational flexibility, there are some things that we want to have for it. We want it to be what we call firm designable. We want to be able to say how much storage goes where. And if you don't have the ability to specifically assign that, then you can find yourself in a situation where you don't have it where you need it.

We would like it to be firm dispatchable. That means we'd like to be able to count on knowing exactly how much there's going to be and be able to make it do what we need to do without worrying about whether it's optionally there or not. It needs to be securable for the same reason as everything else in the grid now because of cyber security and physical security issues. And its service must be assured. We need to be able to count on it and not have it be there at the whim of a business model that says sometimes it's there and sometimes it's not.

So we look at that and say a lot of that points to the utilities being able to control the embedded storage. Now, when I talk about embedded storage, that does not mean that I am saying you shouldn't look at storage that's attached at the edge for various purposes. There are very legitimate reasons to do those things. What I'm saying is there's an additional use of storage that's interior to the grid that gives us this sponginess and springiness that makes it able to deal with all these volatilities that are hitting the grid increasingly. So next slide.

MR. MCDOWELL: Jeff --

DR. TAFT: Yes.

MR. MCDOWELL: -- let me ask you one question from this slide.

DR. TAFT: Yeah.

MR. MCDOWELL: Firm designable, the idea that the utilities should be able to kind of dictate where that storage is to get maximum value out of it, I guess, in states that have taken a position on storage, either through the legislature or otherwise, and putting storage in place, are they going through a very intentional process to say yes, but don't put it all right here; put it in certain locations driven by certain design parameters?

DR. TAFT: It varies from location to location. In some states they've simply mandated that we need a certain amount. Somebody else, you all figure out where you're going to put it. In other places they're trying to be a little bit more deliberate about that, but in no case that I know of has anybody thought about this in a very systemic fashion.

There are places where people are proposing to do that, but in some states there has been concern about the use of storage where there are centralized wholesale markets, which is not the case here, and whether the utility would use that to bid into those markets and be able to have an advantage over other third parties. That's a resolvable issue, but they spent time talking about that more so than saying where will the storage be.

Now, in the state of Hawaii, what HECO has done, and the Commission has agreed with for DER in general, not just storage, but general distributed energy resources, is to say all right, well, if you are a third-party owner and you want to connect, then we, the utility, will -- there's a tariff for doing that, but we, the utility, will control the operation of that device within certain ranges that protect their interest so that you don't have people saying, well, I'm going to have this differentiated set of services and I'm going to charge different amounts for different services. The utility controls it and says I need it right now to do this for me, I need it right now to do that for me. And so the question of ownership is a little bit separate from the question of operational control, the issue being that the organization that knows what the grid needs at any given time for that sponginess is the people who operate the grid.

MR. MCDOWELL: Yeah.

DR. TAFT: Right? So we've seen some states like Texas where the utilities have proposed to be able to deploy storage throughout their grid, and in the case of Texas, the administrative law judge actually went through the arguments from one of the utilities, AEP, and said yeah, good idea, and the Commission said, no, don't want to do it.

So there's a lot of ins and outs to the way people think about it that is very much in flux right now, but we're starting to see more and more people thinking about this springiness/sponginess idea because of the focus on resilience more than anything else. Remember, I showed you that slide and I circled those two things, resilience and operational flexibility? If your objectives in terms of what you want to do with the grid focus a lot on that, this becomes a more important idea. That's not true in every state by any means, so you see a lot of variability in how people are thinking about this. The idea of this as core infrastructure is not brand new, but it's relatively fresh thinking. So in a lot of places for a while now the view has been, well, storage is owned by third parties and it's sold as a service, you know, and it may be behind the meter or it may be attached in some way like an ancillary services device. The idea that it needs to be this sponginess that's built into the grid is relatively new, but we've seen a lot of people thinking about this.

So it's in the early stages of thought, and the reason that I wanted to come and talk about this today is so that you would have this concept, along with all the other ones that you're considering, when you think about storage because my view is the grid as a whole needs to have this capability. It's the fundamental thing that's missing in that complex system that all of our other complex systems have.

This buffering capability is just not there, and yet we are subjecting our grids to more and more volatility from both directions, and that volatility can flow from the edge from the distribution level up into the bulk power systems and create difficulties there. It can flow the other way because we have the sources of volatility at both ends, so to speak, when we have large solar facilities or wind facilities that come from the bulk system and impact distribution. When we have a lot of distribution connected resources, it can go the other way and actually impact the operations of the balancing authority. And we're seeing people start to be concerned about the export of volatility from the distribution level into the bulk energy system. It creates operational difficulties there.

In addition, if people want to be able to do things like use the distribution system as if it is a network for energy transactions, the concept that you see in some states, you know, peer-to-peer energy transactions and all that, then you have to think about can the distribution system actually support that kind of capability, and the answer is with our traditional distribution systems not very well. What's missing is the ability to manage those flows and manage the time differentials involved, and that's the very thing that storage gives you the flexibility to do.

So if the model is that the distribution system becomes a network that facilitates energy transactions, then it's going to become more important to have that storage capability built into it because otherwise you don't have enough flexibility to be able to accommodate all the transactions that people are going to want to be able to do.

MR. MCDOWELL: Good. Thank you.

DR. TAFT: Okay. Okay. So where would you put storage if you're going to use it in this manner? If it's going to be flexibility and if it's going to be buffering for the grid, where do you put it?

So we did some studies about that, and what we concluded was that where you would put it is in the transmission distribution interphase substations on the low side, on the distribution side of those -- on the lower voltage side, in other words, of those substations. That came about through doing some simulation studies and so on and some engineering considerations about what would be most effective in terms of the ability for it to provide that sponginess and also manage that at a reasonable cost. If you have to connect it to the high voltage side, it's a lot more expensive to do than if you connect to the low voltage side. The simulation studies show that connected throughout the system at the low voltage side works very well in a variety of cases that we studied in terms of flexibility. So there’s a rational way to think about where you’re going to put this stuff. Next slide, please.

In terms of operating it, you could treat each one of those as a separate device and treat it as a standalone device, but that’s probably not the most effective way. The most effective way, we think, would be to treat them as a coordinated group of units and operate them collectively. That, again, points to a method of operation that probably works well if it’s handled through either the balancing authority or the actual utilities themselves because, again, the way that you’re going to want to do that depends a lot on the state of the grid, and that information is in the hands of those operators, not in the hands of, say, third parties or even the generators.

We also know from this work that you don’t have to put storage everywhere. You can share it across multiple substations, and we’ve demonstrated how that can work so that you can roll this out incrementally, you don’t have to go out and say, well, every single substation is going to have to have a storage unit. And the other thing that you don't want to do is try to create just a few really big ones and put them somewhere in the system. That turns out to be not very effective. It's also massively expensive. So you can do this incrementally, and you can get the benefits of this that build up over time and do this as a rollout instead of saying, well, I've got to do it all at once to get something useful out of it. Next slide, please.

I mentioned --

COMMISSIONER DUFFLEY: Dr. Taft, could I interrupt --

DR. TAFT: Sure.

COMMISSIONER DUFFLEY: -- just for a second?

DR. TAFT: Sure.

COMMISSIONER DUFFLEY: It's going back to who is operating these storage facilities. What did the Texas judge agree with and the PUC say no to? You referred to that. Was it --

DR. TAFT: Yeah. That was -- that --

COMMISSIONER DUFFLEY: -- that the utilities operate --

DR. TAFT: That was AEP.

COMMISSIONER DUFFLEY: Uh-huh.

DR. TAFT: They wanted permission to put

storage units into their substations and to directly control them for purposes of approving resilience, and they had a list of things that they wanted to do specifically. If you look back at the storage buffer function slide that I had, you'd find them on there. And the administrative law judge looked at their argument about that and said that it was reasonable and recommended to the Commission that that should be something they should be allowed to do. The Commission declined to allow AEP to do that.

COMMISSIONER DUFFLEY: And for what purpose? Was it cost or some other reason?

DR. TAFT: I believe they were still concerned about whether the storage would be used to bid into markets, because they do have markets --

COMMISSIONER DUFFLEY: Uh-huh.

DR. TAFT: -- in Texas, and whether that was going to be fair to other stakeholders and -- as if the utility would have an unfair advantage in operating that bidding into the market.

I will tell you that I don't think that that has to be an issue. We went through this discussion in Ohio about that, and it seemed clear from that discussion that you could delimit the functionality that was allowed and, you know, if you had markets, which you don't have here, you would say, look, they're not allowed to bid those services into the market. This is pretty clear. You don't do that. You use this for things like black start. You use it for things like managing congestion. You use this for things like ride through on outages and so on as opposed to saying, well, I'm going to sell ancillary services. And so you could delimit that, but in Texas they weren't willing to consider that. They just said, look, we don't think we want to go there.

COMMISSIONER DUFFLEY: Okay. Thank you.

DR. TAFT: All right. I mentioned early on in my presentation that because of the convergence of natural gas and electricity and gas being used for generation that there were some interesting opportunities related to storage there.

Gas systems have storage. They have big storage tanks and they can also store gas right in the pipelines by doing what's known as line packing. And you all are probably familiar with that. Basically, there are times when they pump up the pressure to have more gas available there when they seen an issue coming.

That takes a little bit of time and a little bit of a look ahead to be able to do. The gas systems would really like to have relatively constant flow to their loads. Electric systems, as you know, aren't quite that constant. We have daily cycles. We have seasonal cycles. We have all that stuff that goes on. That's why we have peaking generators and that's why we have reserves and we do all that fairly complex stuff to do load following, so there's somewhat of a mismatch in those things.

Well, if we looked at both the storage on the gas side and if we had storage on the electric side of type I'm talking about, you would then have the opportunity to use those two things to even out that mismatch in volatility, too. So it's not just volatility interior to the electric system coming from the various kinds of generation; it's also the connection to gas systems that you would look at storage and say this can help us make that work better as well. And that's what I meant by that originally. You would have to have some reasonable amount of storage on the electric side that we don't have yet today, as well as the gas side, and you would have to have a certain amount of cross observability and coordination.

That cross observability and coordination is being developed, and you may remember that there were some FERC rulings about that after the problems up in the Northeast a few years ago in terms of synchronizing markets and in the terms of literally what’s called cross observability, in other words, sharing state information across those two systems.

Well, those are the basis for being able to do that, so once you can do that, if you have storage on both sides, you have the opportunity to co-optimize the use of that to make those two systems work better.

Same thing when you have other kinds of generation. It’s evening out the volatilities. And that’s what storage really does when you think of it as a shock absorber, is it decouples those volatilities so that the variation of one site doesn’t impact negatively the operation of the other site. In the case of the grid, that goes in both directions, as I mentioned. Okay. Next slide.

So because grids lack this common capability that we have in every other kind of system, one of the things that we suggest to the people is think about the need for internal buffering in the grid. This is a systemic issue. This is not a point issue.

You know, a lot of times what we see people doing with storage is addressing point issues like should

I use storage or should I build another transmission line for reliability purposes and doing the tradeoff between the two, and that's all fine, but what we're talking about here is thinking about the grid as a whole system, which is what we do with the architecture work, and asking ourselves do we want to improve the resilience and the operational flexibility across the board. Do we want to make it possible to deal with these large scale changes that are happening to our grid in general? And if so, perhaps embedded storage embedded in the grid as core infrastructure is the way to go. And if you think about that, you're making a transformation on the grid, giving it a capability you didn't have before, which is this buffering capability. It's a key aspect of resilience in complex systems, and I -- and we're missing it in the grid.

So if you have a focus on resilience, that's why you would be maybe more concerned about this. We know that in quite a few parts of the country that is the case, that there was a focus on resilience. It plays out in different ways in different parts of the country depending on what the vulnerabilities are. And certainly the folks at DOE have a big focus on resilience. In fact, I told people that resilience was the 2019 utility word of the year. In 2018 the word of the year was platform, in case you hadn't seen that.

So there are some key requirements if you're going to do that, where you put the storage, how much of it you use, how it's operated. Those are kind of the things that are key to think about there, and those would influence how you think about how this all gets done in terms of what you as a Commission do in terms of what the utilities would do, in terms of what other people would do.

So if you decide that's the direction that you think is valuable, then there are some recommendations there for how you would actually do it from an architectural standpoint.

And I'm going to stop there at that point and see if you all have some questions for me about all this.

COMMISSIONER DUFFLEY: So going back to -- well, they're not numbered, but it's this Architectural Issues Operation where you show a picture of maybe a storage device on every substation, but then you mentioned you don't have to have this storage, this embedded storage on every substation. Two questions. Like what is the size of these embedded devices, and then what percentage of the sub -- you said not 100 percent of the substations, but 50 percent, 25 percent?

DR. TAFT: So those are actually engineering issues that we typically don't go that deeply into at the architectural level. What you would want to do there is do the engineering studies to determine just how much you're trying to improve that resilience or operational flexibility, and then that would tell you how much storage capability you need.

The work that we did, the simulation studies that we did said that you could look at this in terms of the peak loads on those substations, and what we were looking at is storage that would over a 24-hour period store enough energy for a few percent of that total energy, so it's not actually that large. When you look at what some people are doing in some jurisdictions, they're talking about building these enormous storage units and they talk about things like, you know, being able to run loads for two weeks if the grid is out and so on. I think that's way out of scale here.

What we're talking about here is modest size storage. In terms of how many substations, it's only the transmission distribution interface substations. It's not the regular transmission substations and it's not just ordinary -- like you wouldn't do it in 4 kV, you know, distribution substations. So it's a modest number; it's not a large number. And we also know from our work that you have the option to be able to share storage across multiple substation service areas. So that means putting a storage unit in one substation which supplies the resilience necessary for a couple of substation service areas.

So there are a number of engineering tradeoffs that you can make there, and there's no simple answer to what the exact number is. But our thinking about this is that this is not nearly the scary gigantically expensive thing that you might think that it sounds like from the beginning at all because it isn't that much storage that's needed. We're talking about, you know, a small percentage of the total power flow being buffered by all of this, not the entire gigantic amount of it. So, again, it just is not that large.

Being able to do it this way means you can also do it incrementally, so, you know, none of this stuff that we do with utilities typically gets done overnight, so you do rollouts over a period of years. And by doing this in this distributed fashion, it's very amenable to doing that kind of a rollout over time as opposed to saying I've got to do all of it before I get any benefit.

You get benefit from each piece of it as it adds in.

CHAIR MITCHELL: Just one follow up for you. When did you all do these simulation studies?

DR. TAFT: We've been doing this work over the last three years for the Department of Energy.

CHAIR MITCHELL: And are they -- did you all -- have you all published anything about them?

DR. TAFT: We've published the architectural specifications. The simulation results we haven't published yet because we're still doing some.

CHAIR MITCHELL: Okay. Commissioner Clodfelter.

COMMISSIONER CLODFELTER: Listening to you and thinking about this is very stimulating, thank you, but would it be a fair inference for me to draw that if I were to permit at a policy level or regulatory level, if I just permit unrestrained addition of storage resources at the grid edge all around the grid uncontrolled, unmanaged by the grid operator, that I'm actually maybe increasing the risk of volatility problems on the grid?

DR. TAFT: Could you be a little --

COMMISSIONER CLODFELTER: Well, I don't --

DR. TAFT: -- clear about what you mean about "around the edge"?

COMMISSIONER CLODFELTER: Well, I -- I'm --

DR. TAFT: Are you talking about like behind the meter or a third-party owner?

COMMISSIONER CLODFELTER: Third-party owned. I mean, I'm really asking in the context here of the sort of environment in which we're operating here, in which we've got an awful lot of third-party owned generation --

DR. TAFT: Yeah.

COMMISSIONER CLODFELTER: -- most of it at the distribution level, some transmission connected --

DR. TAFT: Yeah.

COMMISSIONER CLODFELTER: -- and everybody is clamoring to add storage to all of that.

DR. TAFT: Uh-huh.

COMMISSIONER CLODFELTER: And I'm thinking, well, whoa, suppose I allow that and none of that storage is under the control of the grid operator, do I increase my risks of volatility?

DR. TAFT: I don't think you would say that it increased your risk of volatility. I think that there are two things that happen there. One is that just adding it without any sort of coordinated approach to where it is and how it's operated doesn't necessarily get you a benefit. And so in that sense it's potentially a stranded investment. That's one thing to consider.

The other thing to consider is that when you have things like this that are behind the meter and they make the apparent load look different than the sort of real load, it becomes difficult for the balancing authority or the system operator to know how to manage their reserves because they can't actually see what's going to happen. And if there are sudden -- there's a sudden reason why this stuff becomes unavailable and there's a reason how that happens, they can get hit with a sudden shock to the system because they don't know what's actually going on with that stuff.

So why would that happen? Well, we have that problem with solar inverters, and that is that the way the standard was set up for inverters was if there is a voltage fluctuation, they were all supposed to pull off the grid. Well, if you've got all that generation and, likewise, if you have storage that's supplying it to the grid and it suddenly disappears on you like that, you know, that's a big step change. It's a problem for the balancing authority to deal with.

If they don't know how much there is because it's in third-party hands and they don't know what's being supported by storage and what's actual variability, they don't have that visibility, and a lot of system operators have been concerned about not understanding what's going on in that combination because they don't have the observability, they don't have the measurements to tell what's going on. So some people's answer to that is, well, you know what, we'll put extra metering in so we can see that piece, the DER piece, separately from the traditional load piece and get more visibility. You know, there's a lot of different ways you can play this out.

So I think the answer to your question is I don't know that it creates additional variability. I think the problem is that you may not get what you were hoping to get from it in terms of resilience and flexibility for operating the system, and you may cause some problems with sort of disguising the actual load because there's no way to tell exactly what part is storage and what part is real load.

COMMISSIONER HUGHES: You mentioned in your comments that you're agnostic about business models, and I'm just trying to wrap my head around -- could you give an example of business models that would fit with all of your simulations, because it seems to me that business models are embedded into your analysis. But could -- if you give me an example, it might help.

DR. TAFT: So I'll give you an example from a slightly different perspective. Some of the folks who aggregate distributed energy resources have argued that in those places where there are organized wholesale markets and if there are going to be distribution level markets because, you know, that's considered in some places, they want to be able to bid into both and they want to be able to be unrestrained into how they go to both, and that has led to this question of how you do transmission distribution coordination in the presence of DER, which is what the distribution system operator conversation is largely about, right? So if your business model is I should have unrestrained access to both markets whenever I want and nobody else has anything to say about it, that's the kind of issue we try to be agnostic about. We don't try to say, well, you should or should not be able to have access to markets.

What we look at is to say what architectural structures will enable people to do what they want, and we're not here to say there shouldn't be aggregators or there should be aggregators or they should have this role or that role. Our argument is what structures enable people things to do -- do the things they want to do and where are the legacy constraints that need to be relieved that would prevent them from being able to do them. So that's what I meant about business models. We don't try to advocate for or advocate against any particular way that somebody might choose to be able to be compensated or make money off of any particular technology related to the grid.

Now, what I did say is that I thought that storage that's embedded needs to be controlled by the utility. That's because they have the state information to know what to do. But I didn't say anything about who owns the storage. Did I answer your question?

COMMISSIONER HUGHES: Yeah. That last sentence answered my question. Thanks.

CHAIR MITCHELL: Any additional questions? Thank you. Oh, Kim.

MS. JONES: Thank you. So I'm having a hard time getting my head wrapped around the benefit of resilience, so where I'm starting from is thinking of resilience in terms of having fewer outages to customers, or when they do happen, you're able to get the lights back on more quickly. Help me understand how having this device in a substation will preclude outages or help you bounce back faster.

DR. TAFT: So first I will say the way that we have defined resilience separates out recovery from the first part. So for us, resilience is largely about not having the outage. And if you have an outage, that is now in what we would say is in the reliability domain, because when you look at the reliability metrics, you're measuring fundamentally two things, how often things -- how often power outages occur and how wide they are and how long it takes to recover.

So an example of how you would improve the resilience, let's say that you have storage in the substation. Let's say that you lost power from your bulk power system and the storage helps you ride through that for your loads. So they don't see the outage even though there was -- there would have been an outage if the storage hadn't been there. So that's a simple example, is ride through on outages.

MS. JONES: So but if -- what if that doesn't ever happen? I mean, that's the kind of outage that just doesn't happen.

DR. TAFT: Well, when you look at resilience, so this -- I was going to spend more time on this because this gets really interesting.

The way that resilience is talked about in the industry is it has a lot of weaknesses to it. And so we didn't define it, and you and I started talking about it, and that's how we get into this sort of mismatch here. Resilience, the way we look at it, is a combination of vulnerability and how you deal with that vulnerability to either prevent an outage from occurring or minimize the extent of it when things start to go bad.

You can't predict when these events that people are concerned with are going to happen unless you get pretty close to them. I mean, when you see -- you can see a hurricane coming when it's finally coming, but is there going to be one this year? Is there going to be one next year? Nobody knows how to do that. So people for a long time were talking about resilience in terms of it being the ability to deal with large scale, but rare, events. And it gets to exactly what you just said, what if it doesn't happen.

All right. So do you have insurance on your house? What if your house doesn't burn down? It's a little bit like that. We know that there are various kinds of vulnerabilities and various kinds of threats. We can't predict when they're going to happen, but we know we have vulnerabilities to them. So the way that we view that is that you need to think about that in terms of the vulnerability which exists today even if the external event doesn't happen today.

One of the things that happens with reliability calculations in the utility industry is it's -- in a sense it's weird because in other industries like electronics and aerospace, reliability is a forward-looking view and it's not based on conflation with external events. In the utility industry reliability is a backwards-looking view and it's based on conflation with external events. So how do we calculate those things? We look at the outages that happened and we calculate all the statistics and we try to figure out from that what to do.

But if you look at -- say, in electronics we look at the characteristics of the components and we look at the structure of how they go together, and then we ask ourselves the question what is the probability of zero failures in a particular period of time, and that's how they think about reliability, not how we think about it in this industry.

So you have that problem that they were conflating, you know, certain aspects of resilience with reliability. So when you say resilience is about large, rare events, you're really saying resilience is a special case of reliability, and what we found is that people didn't get good answers for what to do when they would get tangled up with that. So we said, no, what you have to do is think about resilience is the stuff that happens before an outage. It's how you resist these problems and keep things going, how well you keep them going, but once an outage occurs, that's now in the reliability domain.

So think about that as the vulnerability exists all the time even if the external event is unpredictable. And you -- and the question is do you want to be ready for that external event, not knowing when it's going to happen.

MS. JONES: Just a real quick follow up. In this model of putting storage at substations and using it as this buffer, in that kind of a world would the utility be able to have a reduced reserve margin from what it would otherwise have?

DR. TAFT: I would say if they have thought about the engineering of that and decided how much storage they would need, yeah, they would be able to trade that off against reserve margin, yeah, but it is a careful engineering calculation to do that because of the need to assure service.

CHAIR MITCHELL: Thank you, Dr. Taft. Okay.

Any additional questions?

(No response.)

CHAIR MITCHELL: I think we will move on. Dr. Mills.

DR. MILLS: Okay. Thank you for the opportunity to speak here today. I'm going to be talking about three different topics that are not necessarily completely tied together, and so there's going to be some gaps between them, but I'm happy to take questions as we go along just to try to fill in some of those.

And so these three topics are first looking at the contribution of solar to overall resource adequacy needs and then the role of storage in increasing that contribution. The second one is to look at some of the literature that's been out there on solar integration costs, and this connects to some of the questions that you all are dealing with now and sort of what drives those integration costs. And then the third part is to look at using storage to reduce solar variability through sort of a mechanism that we refer to as ramp control, and compare the cost of doing that with storage, that ramp control with storage to these integration costs, and that's where it gets a little bit loose. We haven't done a lot of detailed comparison, but kind of give you some order of magnitude estimates from the work that we've done.

And much of this work that we've been doing recently is based on some work in Florida, that we've been working with some of the municipal utilities there, and they've had a lot of -- these questions have been driving the analysis that we've been doing. This was all funded by the Department of Energy through the Solar Energy Technology's office. And then also we've done a fair amount of technical assistance to states that's been funded by the Department of Energy Office of Electricity, and so I'm going to be pulling from different parts of that. And, again, it's not really a comprehensive analysis of all of these different value streams, but more it's sort of some of the methods and some of the insights that we get into some of the various key elements of it. Next slide, please.

So first I'll jump into looking at the solar and storage part of this for resource adequacy. Next slide.

The main part of this that we looked at was trying to understand how adding solar can increase sort of the ability of our system to reduce peak demand and then how storage can help increase that benefit. And so the idea here is to be looking out with this red line sort of as your utility forecast over time that says here’s what our peak demand might look like over time, and in order to meet that peak demand, we might have to build plant in a particular period. And then if we add PV or we add PV and storage, we’re going to have this new trajectory that’s going to be our peak demand now being lowered because of the addition of that asset, and that might defer the need to build this capacity resource some years.

So the ability of solar or solar and storage to kind of create a gap between those two lines, between the peak demand without PV and the peak demand, the red and blue lines, that ability is really driven by what we call the capacity credit of solar. So that’s sort of the fraction of the nameplate capacity that contributes to lowering peak demand. And that ability to defer the need to invest in that power plant, when you sort of look at the net present value of that, that becomes the capacity value and what we refer to as capacity value in dollar terms. So this is really the economic value of that capacity credit that you’re getting. And so our analysis, we’re just going to focus in looking at this capacity credit and seeing how that varies with different factors. Next slide, please.

Really, this is not something where we were trying to go in there and answer this from a detailed perspective and sort of get the right number; instead, what we were trying to do is to develop some simple methods to look at all the different factors that might affect this capacity credit and then pull on a bunch of them, try to find out what are the things that really drive this calculation and what are some of the factors that you should be aware of when you do then go into a more detailed model.

Some of these detailed models really are expensive in the sense that you have to set up a lot of assumptions and parameters to them, and so you have kind of few chances to really investigate a lot of different directions. We wanted to come up with a method for simplifying this and then just get some intuition out of it.

And in part, the work that we were doing was in parallel with the National Renewable Energy Lab, was using their resource planning model, which is sort of like an integrated planning model, that looks at some of these capacity credit analysis internal to the model. So we're trying to sort of unpack a little bit of that and get some intuition for why you might get some of the results out of that capacity expansion model.

And, again, what we're really looking at is what are some of the factors that drive these relative changes in the capacity credit, and we're not trying to get a very precise estimate of it for one particular configuration. And the idea was that that would help kind of prioritize additional research directions and questions and sort of see where it might be most interesting to dig in further. Next slide.

So with this we did start to take solar and add storage to it or we looked at storage independently and looked at the capacity credit of storage, too. And so one of the things just to note with this is that one of the things we really wanted to vary was what we refer to as the duration of the storage. And you might be familiar with sort of storage being rated in terms of its ability to -- the rate of power it can charge or discharge at, so in this illustration that's at 10 MW and the amount of energy that it can store in there. In this case it's a 40 MWh battery.

And so the duration of it is what we refer to as how fast that reservoir would be drained if we were to discharge at full capacity. And so in this case that 40

MWh reservoir, if it was full, if we discharged at full capacity it would be drained in four hours. So that’s a four-hour duration battery, and we’re really -- one of the things that we’re going to look at quite a bit is how the capacity credit of storage changes as a function of that duration.

One of the things to note here, too, is that we were, again, simplifying a lot of things in here, and so we’re going to treat that battery as fully chargeable and dischargeable when in reali--- so in reality, people will oftentimes restrict how much of that reservoir they actually access to preserve that asset life, and so our sort of reference to four hours is sort of meaning the accessible energy that could be in there which might be different than the true rated capacity of that. And, again, that just comes from, you know, operation people might hold some of that back to avoid degradation. Okay. Next slide.

So in order for us to understand and just get some intuition as to what the contribution of solar and storage is to meeting our resource adequacy needs, we started off with this simple idea of just taking a load duration curve. And so the green line on the image on the left takes the load in this particular utility and sorts it over every hour of the year, so 8,760 hours, from the highest load to lowest load level. And in order to make sure that your system is adequate, you're going to have to be looking at some of those peak load hours. Those highest load hours are the ones -- really going to drive the need for having adequate assets on the system to meet those peak demand needs.

So then our question is, is if we add some asset like adding solar or adding storage, we're going to reduce that load in certain hours, and in particular what we want to focus on is how much can we reduce the load in those peak hours. And our ability to reduce load in those peak hours is what's going to sort of allow us to avoid the need to build other assets to meet that peak demand.

And so in this case the chart on the right zooms in to just the top 100 peak hours of the entire year, and our green line, again, is that load sorted from highest load to lowest load hours. And then the orange line, we've done the same thing now, but it's the load minus solar in this case. And then we sort it from highest to lowest. Or we could do that with load minus the storage generation and then sort it from highest to lowest.

And so the gap that emerges between those two, between the green line and the orange line, that becomes our capacity credit. That's our estimate that -- the average amount of gap between those is our estimate of the capacity credit of that asset.

And in our work, one of the things that we focused on was trying to come up with a fairly fast and robust way that we could dispatch storage, find the way that storage would be dispatched in order to maximize that capacity credit. So we sort of were looking at an idealized situation where you had full control of that storage and you were able to dispatch it within its capabilities such that it could maximize that capacity credit. And just to explain the chart a little bit, that orange area that's filled in, if we could minimize that area, that would be the same thing as maximizing that capacity credit. And so a fair amount of our work went into developing methods that would be fairly quick to calculate that. Okay. Next slide.

So then one of the things that we started to do was just to kind of parse this problem out of how PV and storage would contribute to capacity, to resource adequacy, and we started by just looking at PV alone, and we looked at various sites around Florida and also various utilities, and looked at how the capacity credit might just vary from different site and utility combinations. And then we also looked at how that capacity credit might change as we deploy more and more solar, to look at sort of any effect that increasing penetrations of solar would have on its ability to continue to contribute to peak demand needs. And we did some questions just focusing on storage alone.

So, again, one of the things that we were really interested in is looking how the capacity credit changes as a function of its duration. If you have a bigger and bigger reservoir, does that allow you to get more and more capacity credit and do you hit any sort of limits on that.

And then another question that we asked was if we start to add a lot of storage to the system, does that change its contribution to resource adequacy. Same thing as what we were asking for PV, is there sort of any effect of penetration on the capacity credit of storage.

And then finally, we looked at combining PV and storage together, and we looked at a variety of ways that you could configure the PV and storage together and saw how that would affect the capacity credit, and then also looked how that might change depending on how you size the battery relative to PV.

So I'll kind of go through some of these results and, again, please do jump in if I can clarify anything, please.

So the chart on the left here shows the capacity credit as a function of increasing amount of deployment of that PV. So as we move to the right on that chart, you have increasing amounts of solar deployment, and the capacity credit is shown on the vertical axis.

The three different lines here represent three different utilities in this area, so we were working with data from the City of Tallahassee, which is a very small utility, only about 600 MW or so, the Jacksonville utility, and then Florida Municipal Power Pool which includes both Orlando and some other smaller utilities.

MR. MCDOWELL: Kim, you need to advance the slide.

DR. MILLS: I'm sorry. Please advance the slide on that. I'm doing this in two places here.

So, again, that was the chart on the left that has these -- these three different lines, and the top one is the Florida Municipal Power Pool, and then the bottom two are Tallahassee and JEA.

So one of the things that stands out is that there are differences in the capacity contribution of solar, and one of the things that's going to drive that is how well your solar deploy--- your solar generation profile is aligned with that time of highest peak demand. And we see that in FMPP which is, again, sort of more in central Florida, the contribution of solar to meeting that peak demand is somewhat higher and then it declines as you add more and more solar.

In contrast, for JEA and City of Tallahassee the capacity credit is quite a bit lower there, and part of that has to do with the fact that both of those utilities also have some of their peak hours occurring in winter and sometimes even at night, and so this is maybe a little bit closer to what you might expect in North Carolina where you do have some increasing amounts of winter loads driving the peak, is that because solar is not going to be -- its production is not as well lined up with those peak hours, the capacity credit will be somewhat less than what we saw with FMPP.

And then to explain why that capacity credit declines as you go to increasing penetration, but even in places like FMPP where you might be aligned with, say, the summer peak hours, as you add more and more solar, the time of that residual assistant peak is going to shift away from solar hours and into hours into the early evening when there is no sun. The sun goes down and the net load is still there. And so that means that the next increment of solar that I add is not going to be able to reduce that peak demand by as much. And so we see this declining capacity credit with solar as we go to higher penetration levels. And that happens in both of these two different climates.

So then on the chart on the right is where we’re looking at just storage alone. And now, again, the vertical axis is looking at the capacity credit of storage and then the horizontal axis is increasing amounts of duration.

So the first thing that is kind of surprising with this is that for something that’s a perfectly dispatchable asset, you don’t immediately just say, well, let’s give it a hundred percent of its nameplate rating, that the capacity credit of storage can actually be quite a bit lower than 100 percent when you have short durations of storage. So short duration storage, something down in, say, two hours or so might only achieve about 50 percent of its nameplate capacity, contributing to you being able to lower your peak demand needs.

As you increase the duration of that storage and that it can store more and more energy and you move out towards four or five hours, you do start to achieve closer and closer to a hundred percent of its nameplate capacity. And that was true for all of these three different load profiles that we looked at.

Now, the difference between the green line and the red lines here, the red lines being on the bottom, is that the red lines are what the capacity credit is of storage if we have a lot of storage on the system. So on the green lines that was just sort of the first increment of storage, and the red lines were meeting about 20 percent of your peak demand from storage. And so you can see there that as you have more and more storage in the system, you require longer and longer duration in order to achieve that full 100 percent capacity credit. And it might go out as far as nine or 10 hours of storage duration before you're able to achieve that.

Well, just these next charts -- if you go to the next slide -- just sort of illustrate that a little bit, that this is one of those utilities where we're showing the daily load profile on a peak day in the winter on the left and then in the summer on the right.

And we have that peak load being reduced by the storage in the dash lines. And if we have just one hour of storage, we can clip off a little bit of the very high peaks, but you're not able to do much with that. And as you go to fours or six hours, you're starting to be able to clip off a larger portion of that peak.

But you can think about it that as you add more and more storage to this, the sort of residual peak that you have to clip off becomes wider and wider, and so that's sort of where you get this idea that in order to get that full hundred percent capacity credit, you're going to have to move down a wider and wider peak that you have to reduce, and you need longer duration storage in order to do that.

Okay. Next slide. So one other way to think about that is instead of saying how long does the duration have to be in order to get that hundred percent capacity credit, you could also look at this as what is the capacity credit of just a fixed duration. So if we just had a four-hour battery and we look at what its capacity credit would be as we added more and more storage, if we start here in the green lines, that when you're at -- when you're at sort of the first increment of storage being added to the system, you can achieve that sort of 80 to 90 percent capacity credit. And then that four-hour battery, as you add more and more storage, will only be able to achieve about something like 50 percent or so as you've gone out to about 20 percent of your peak demand being met by storage. So that's that idea again, the peak gets wider and wider.

Now, the red lines in this case that are above that are asking a question of how does that capacity credit storage change if we actually have a lot of solar in the system. So one of the things that solar does, is that it can kind of narrow that residual peak demand that's left over. So particularly you take somewhere like FMPP where, again, we have peaking that's happening in the summer, well, with no -- with no solar in the system, we have a pretty wide daily peak that has to be met by storage, and so it takes four or five hours of storage to be able to meet that peak demand.

Now, as we add solar into that system, it's going to reduce the peak demand during the day and shift that more and more into the night, and it becomes sort of what we refer to as a skinny peak. So it's a skinny net load peak that is left over in the night because of that solar, and then that means that the storage and what it has to do in order to continue to meet resource adequacy needs becomes somewhat easier. And so you can get a higher capacity credit for that four hours of storage. Again, though, as we increase the amount of storage, that you still do see this decline over time that can happen, but that explains why the red line is higher than the green line. Okay. Next slide.

So that was sort of looking at the storage and the PV as somewhat independent systems. And then this final part of it we wanted to look at what would happen if we start to couple the PV and solar together and what effect would that have on the capacity credit of that combined asset.

So when I refer to it as independent, that's, again, that's just this sort of you have a PV system over here and a storage system over here, and they're not sharing any equipment and you're not restricting the storage to charge from the PV or anything like that, so they're independent systems.

One thing we could do is we could just then bring those two together where we might have the battery sharing the inverter with the PV system, so it's behind the inverters. We call that DC coupled. But we still allow that battery to charge either from the grid or from the solar. So this is what we call a loosely coupled system. So it’s sharing some of the infrastructure, it’s sharing that inverter, but it can charge either from solar or from the grid.

And then we can go to a tightly coupled condition where it’s, again, sharing the inverter. It’s on the DC side of that. And we’re restricting that battery to only charge from the solar. And this is sort of the case that you get towards because of the investment tax credit which does actually require that for a battery to get a cost reduction or the tax credit applied to its capital cost, it has to demonstrate that it's being primarily charged from the solar asset. So the tightly coupled case is sort of starting to get more and more towards what you might be pushed for because of the current tax policy. Okay. And then next slide.

So this is now kind of looking at that capacity credit of the PV and storage system. And the chart on the left is FMPP, again, kind of down in the center part of Florida. The chart on the right is JEA. JEA has sort of that more of a high winter peak and some summer peak, whereas FMPP is really dominated by its summer. And so if we were to -- and the horizontal axis on both of these cases is increasing amounts of hours of storage, increasing that storage duration.

So if we start at zero hours of storage duration, then the capacity credit that we get of any of these systems is just whatever you get from the PV itself. So it starts off at the capacity credit of the PV itself, and then as you start to add more and more storage duration, you're now combining the capacity credit of the solar plus whatever capacity credit you would get of the storage asset.

Now, at some point there starts to be some differences in these lines depending on the configuration of it, and that happens around two hours or so. So for the independent system we just kind of keep going up there and we're adding the capacity credit of the storage to the PV by themselves, and that's what the red line shows, so that's what you would get if these two were completely independent, whereas the purple and green lines are showing that you start to run into a limit which is the inverter of that system, so the battery in the PV system are behind a shared inverter, and in this case that inverter is 100 MW, and because of that, that shared inverter you're sort of limited how much total capacity credit you can get by that inverter, and so it sort of caps out at 100 MW. And so by coupling them, we're actually getting a reduced capacity credit.

Now, in FMPP where, again, it's summer dominated, we don't see much of a difference between whether we restrict the PV system to charge only from the solar or if we allow it to charge from the grid. Those two end up being about the same. On the other hand, if we go over to JEA, we do start to see a difference where if you restrict the PV from -- I'm sorry -- the storage from charging only from the PV, that we actually start to see a lower capacity credit. And this, again, has to do with some of those winter events that happen. End up in a situation where you have a winter peak that's coming by, and your storage system would be required to only charge that battery from the solar, and if there hasn't been much solar in that day because of clouds or things like that, you're not going to have sufficient energy to charge up that storage and it won't be able to contribute as much to that winter peak. And so that's that slight difference that we see between the purple and green lines, is sort of that effect of not having sufficient solar energy to charge the storage system.

And, in fact, at some of these lines you can see that that coupled system might even achieve a lower capacity credit than just the storage by itself if the storage was allowed to charge and discharge from the grid. So that's the blue line there, is storage alone.

I think one of the things, you know, to look at is that this helps us build some intuition in some of these story lines and try to understand the interactions of some of these. At the end of the day, these are not huge differences between, say, the tightly coupled and loosely coupled. These are, you know, kind of within the range of some of the uncertainty on that. But it does start to kind of play out some of these issues and lead to things that are worthwhile kind of investigating and thinking about a little bit further where you're considering different opportunities for structuring this.

Okay. And then the final slide -- final couple of slides here -- next slide -- are just to kind of go through that even though we're sort of doing this in a very -- a method meant to kind of explore some of these issues, we did want to try to understand how close this would get to something that's a little bit more realistic.

So in the -- in this chart here we're showing some of these capacity credit results focused either on storage on the left-hand part of the chart or solar on the right-hand chart, and our approximation method that I described is what is shown by the blue bars in this. And so each of those is just reiterating some of the results that you've already seen before. But what we do is then sort of benchmark that against a much more detailed and, again, expensive to run probabilistic approach.

And so this probabilistic approach is what's kind of more of the gold standard in really trying to understand the capacity contribution of different assets, and so this is referred to as the effective load carrying capability, and that's a much more rigorous method for doing this, but it's expensive, and so it's harder to do and to explore some of these issues so we just did this in a couple of cases.

And what you can see is that for our JEA and City of Tallahassee -- oh, sorry -- I'm sorry -- JEA and FMPP, the blue and the green ones are not too different and sort of get some of these trends that look fairly similar to each other. Where we see a huge difference is in the City of Tallahassee, where when we do this probabilistic approach, it's kind of more of the gold standard, we get distinctly lower capacity credits. And this has to do -- this is a very sort of specialized case where it has to do with the fact that the City of Tallahassee has just a couple of very large generators relative to its size. It's only about a 600 MW utility, and it has some generators that are about 200 MW, and the loss of one of those generators can sort of cause you to have a lot of outages. So in the case of the City of Tallahassee you actually have a lot of risk of outages spread out over a huge number of hours, whereas in FMPP and JEA it's much more concentrated in sort of those peak demand hours.

So our approximation method, we're really focusing on sort of the top 100 hours of the year, and that does okay compared to this probabilistic benchmark, but in the City of Tallahassee where that risk is actually kind of much more widely distributed, focusing just on the peak hours doesn't do as well here, and that's what we found. I think in most places you don't have that situation that the City of Tallahassee faces, so we feel pretty comfortable. For most places this approximation method does yield some valuable insights. Next slide.

The other thing is that in all of what we were doing is that we're sort of operating the storage with perfect foresight, so we're basically optimizing the dispatch of the storage, seeing some historical weather year of load data, and so that's going to be optimistic because you can't truly implement that. That's not something that's implementable.

So we looked at sort of creating a -- kind of a bookend to that of what is an implementable case that doesn't require very much information. And so this is what we -- you might refer to as back casting or sort of like a day-ahead persistence, where what we did was we said, okay, we know what happened yesterday. Let's dispatch our storage in an optimal way, given what happened yesterday, and then implement that today without looking at today at all. So it's something you could put into practice, and you could implement that.

So we have our optimistic case and then sort of this pessimistic case, and reality is going to lie somewhere in between those two. And what we find is that if you have longer and longer duration storage, that your sort of perfect approach is going to be pretty achievable with this sort of day-ahead persistence approach, that you can get about 80 percent of that same capacity credit at least with -- if you have long duration storage.

If you have fairly short duration storage and we only have a couple of hours, then it becomes pretty important to dispatch that storage in exactly the right times, and so that makes it so that forecasting becomes much more valuable, and doing this based on what you observed yesterday is not going to achieve the outcomes that you want. And, again, that was particularly true for the City of Tallahassee, where in this worst case if you just had one hour of storage duration and you used sort of yesterday's information, you'd only achieve about 30 percent of the capacity credit results that we showed earlier.

Okay. So that covers it for what we have looked at with this capacity credit part of it. Now what I'll jump to in the next couple of slides is to start zooming in a little bit more on what's happening within the hour. So a lot of what we were describing is sort of something you can do on an hourly basis, and we're looking out over the whole year and we're sort of focusing on these longer-term planning issues.

With variability we're starting to talk about more of what's happening within the hour, and it becomes much more of an operational issue. And a number of studies -- actually, a number of models don't do very well at sort of looking at these operational issues in that short term. And so a lot of studies have done these integration cost estimates to sort of fill in those gaps and say if our models can't really capture that very well, what's missing from it and how much might that short-term variability cost. Next slide.

So we've seen that a number of entities have conducted some of these studies to estimate integration cost. There's a lot more literature out there on wind, and it goes back into the late '90s, early 2000s in the U.S, and then there's a few studies that have been done on solar. There's an increasing number of studies that have been done on solar, but the data I'm going to show you is somewhat dated because we haven't tracked that as closely.

One of the things that you do see, though, is that there can be a huge variation in these integration costs. We do see a lot of variation from study to study. And part of that has to do with the different resource mix that if I'm integrating wind or solar into a system that has a lot of inflexible baseload units, for example, that's going to have a different impact than if I have a system that's got a lot of flexible small combustion turbines that can fire up and help out a lot. So that resource mix matters.

Institutional setting matters a lot. If we have very large balancing authorities that are sort of coordinating over a very large footprint, and we're doing that through things like organized wholesale markets or things like the energy and balance market in the West, that may make it easier to integrate this variability than if you go to a setting like a municipal utility that's running its own balancing authority. So if I have a very small utility with a very small footprint, the integration challenges are going to be much larger in that case. So that institutional setting matters quite a bit.

The final part is that there isn't really a standard way to define these integration costs, and there really isn't a standard methodology for it. There are some best practices that are out there, and some studies kind of follow those to different degrees.

And part of these integration costs, again, is that you have to sort of identify what's the purpose of it. What role is it fulfilling and why do I need this integration cost estimate? And in some contexts the reason why people will do these studies is for integrated resource planning. Again, they might have a capacity expansion model or a production cost model that is going to drive a lot of their decision making about what assets to add and what sort of costs they have, and those models don't do very well at really capturing some of this very short-time scale operational issues. And so integration cost studies can be designed to sort of fill in those particular gaps. And so some of these studies have been designed in that particular way, whereas others have not as much.

And so, again, you just sort of have to phrase this question of what is the purpose of this integration cost and what aspects are maybe missing from sort of what our standard methods are that we're going to be filling in with this integration cost study. Next slide.

So as part of an annual market tracking report, for a number of years we've done a survey where we'll just go out and look to see what sort of integration cost studies have been done over the years, and we'll collect that information and then put it into this chart here. So this has been something that's, I think -- I think we've done this at least through -- since 2007, so it's a number of years, and we just add data points to it as they come along. This, again, focused on wind. And it suffers from all of these limitations that I mentioned before, that there are a lot of different variations and methodologies and institutional settings and resource mixes, but at the end of the day you sort of get this range of integration cost. And then there's different amounts of wind that are being added in each of these studies.

I think one of the things that stands out to us is that a lot of them to tend to cluster in somewhere below $5 a MWh. There's a few outliers that have this extremely high cost that ramps up pretty high. The one that stands out here I believe is from Idaho Power, which is sort of a fairly small balancing authority that has a lot of wind that’s being added to it, and so they sort of face in their studies, they face increasing cost of trying to integrate that wind into their system, and that’s what these costs reflect; whereas in other places, if you go to some of those ones on the very bottom, Southwest Power Pool, for example, is a place that has a lot of wind, but it’s spread out over a fairly large footprint, and the studies that have been done there show a fairly low integration cost below $2 a MWh even for very large amounts of wind. Next slide.

COMMISSIONER DUFFLEY: Mr. Mills?

DR. MILLS: Yes.

COMMISSIONER DUFFLEY: So the Idaho Power example, is that because of transmission upgrades or what’s --

DR. MILLS: No. Oftentimes these are just going to be driven by operational costs. So a lot of times what you might do --

COMMISSIONER DUFFLEY: Just the operation --

DR. MILLS: -- is you might sort of say here would be what my cost of the system would be if I had a certain resource mix and maybe a certain amount of wind that's predictable and not variable within the hour. And then I might look at -- because of the true unpredictability and variability of that wind, I might have to increase my reserves or do these different things on an operational basis that will then impose cost, and you're just isolating those operational costs. So it really comes from startup and shutdown costs, sort of running your power plants at part load, and then any sort of additional capacity cost associated with reserves. Next slide.

So for solar, again, we don't track this as regularly. There was a nice kind of meta-analysis that was done by Synapse a number of years ago, and they grabbed a number of integration cost studies that had been done for solar across the U.S. There are -- again, one of the things that we see is that a lot of these sort of show something in the sub $5 per MWh range, but there's a lot less to look at here.

One of the studies that's listed as APS-Argonne

2013 is one that I was a lead author for, so we worked closely with Arizona Public Service to look at trying to quantify these integration costs in the context of their integrated resource planning. And so we were specifically trying to identify those things that are missing from their production cost model and do a very detailed operational model and then look at the difference in those costs and then quantify that and come up with it. So really this is driven by primarily two factors. One of them is your ability to forecast solar on a day-ahead basis and the fact that you have to make some decisions about which units to turn on and off. And the second is sort of that within the hour variability that caused them to have to hold more reserves, and so we quantify the additional cost of those reserves.

And in our study we had some various sensitivities and ranges on that. The number that Synapse pulled from that was probably our base case where it was somewhere south of $4 per MWh.

COMMISSIONER CLODFELTER: I’m not going to go into the substance of it because we probably shouldn’t do that in the context of this hearing. I’m just -- so this is really just a question of curiosity. Do you have any familiarity with the study that was done for Duke Energy by Astrapé on solar integration cost?

DR. MILLS: Only in sort of preparing for this where I've skimmed some of the documents that have been done about it, yeah. And I do know that one of the things that, you know, they quantify in there is the additional reserve requirements.

COMMISSIONER CLODFELTER: Right.

DR. MILLS: Those didn't stand out as being remarkably different from what we had come up with. And then I think in general, some of their methodology is different than what we had done with ours, and I don't fully understand the implications of some of those differences.

COMMISSIONER CLODFELTER: I didn't mean to get into the substance of it.

DR. MILLS: Okay. Sure.

COMMISSIONER CLODFELTER: I just -- because I don't think we should do that --

DR. MILLS: Okay. Yeah.

COMMISSIONER CLODFELTER: -- here, but just wanted to know if you're familiar with it.

DR. MILLS: Yeah. And I guess just for context, just on that last one I think that the numbers from the Astrapé study are somewhere even below the $2

MWh, they're somewhere in that range or around there, so kind of on the low end of what this is seeing. Okay. Next slide.

I think one of the things that to us jumps out is that these costs are not particularly high. So this is sort of you're quantifying the cost of all of the within hour variability and some of the forecast ability issues of solar, and they don't jump out as being, you know, these huge numbers. You know, if you're talking something that's down in the couple dollars a MWh range, that sometimes can be surprising. And a lot of that does come from the fact that when you -- these studies are being done, oftentimes what you're doing is that you're looking at trying to find those additional reserves that you might require or dispatch of your power plants when you're aggregating all of that solar over the footprint of your balancing authority.

And so in our case when we were doing this with the Arizona Public Service, we were looking over, you know, largely from Yuma to Phoenix or so, which is, you know, on the order of a few hundred miles of distance and a lot of solar plants being distributed throughout that. And so what's happening at an individual plant is that as a cloud passes over, you might see a huge variation in that power plant output, but that variation won't necessarily be correlated with what happens at the neighboring plant and then the one that's a few hundred miles down the road, so as you start to aggregate that, and the power system can act a little bit like a big bathtub, you can tend to smooth out a lot of those variations.

So this chart here is from that study that we did where it shows our modeling of individual power plants, solar PV plants, there were 32 of them in the red lines on the left, and then shows what happened on a fairly clear day relative to a day that was partly cloudy. So on the clear day you can see that for the most part, the power plants come online early morning, have fairly steady input, output during the day, and then drop off at night. You do see some instances where clouds come over and different things were slightly affected, but it's not very large. And so then the aggregate output is fairly smooth in the black lines on the right.

Then the next day we had quite a bit of variability that was happening, and you can see that some of these power plants are jumping between nearly their full output and zero output as those clouds are passing overhead. So we see at an individual plant you're going to see a huge amount of variation that would occur, but as we aggregate that over the balancing authority footprint, that for this amount of solar you see a -- you never see an instance where it's going to jump through that full range. We do see more variability on the partly cloudy day, but it's not nearly as bad as it is at an individual power plant.

And so that means as we're looking at what the costs are from a system perspective to manage that variability, the reserves that we're adding are not reserves based on sort of each individual plant, but instead it's sort of that residual aggregated output. And so you might end up finding that the amount of reserves that you require to manage that sub-hourly variability might only be, say, a few percentage of the PV nameplate capacity. So if we have 1,000 MW of PV, we might need something on the order of 10 or 15 MW of reserves, but it's not 1,000 MW of reserves that we add. And so that's why some of these costs are actually quite a bit lower than you might otherwise expect. It does have to do with that aggregation.

And that sort of sets up for why I'm going to go into this next stuff, where we're going to look at trying to control variability at an individual power plant using storage, and that's because in places like Florida where you have municipal utilities who are their own balancing authorities, they don't have that ability to aggregate over a large footprint. They're looking at an individual city, and so everything is sort of happening very locally and they are running into issues managing that variability at a very local level, so you then turn to solutions at the plant level. In North Carolina, on the other hand, you do have balancing authorities that span nearly the state footprint. Okay. So next slide.

So that's going to tee up this last part where, again, what we're going to be looking at is sort of driven by some of the questions that came up in our work with Florida, was really just to look at if we did start to restrict how much variability we wanted from an individual power plant, what would -- how -- what would it take to do that using storage. And then our question was how much would that cost to do that? How much would you have to pay in terms of storage? And then let's start to compare that to some of these previous integration studies. They're not exactly comparable. They're not apples to apples, necessarily, but in terms of orders of magnitude, we can kind of start to get a little bit of a sense of if doing it locally with storage might have somewhat comparable cost to doing it through aggregating at the balancing authority level. Next slide, please.

So the first thing we needed to do was to size the storage system in order to meet some of these ramps, so a key parameter that we're going to be looking at is a ramp control requirement. And so if we have a -- basically, a ramp control requirement says what's the maximum ramp that a PV plant can go through per minute in terms of the percentage of its nameplate capacity per minute. And that would be a restriction that we're going to apply, whereas if you have something that's 20 percent per minute, that's pretty relaxed. You know, with a completely uncontrolled plant you might see something that would go as high as 50 percent per minute or something like that. So you might go to 20 percent, and then if you want to get really strict, you can go down to, say, like 2 percent per minute. So you're saying the most the PV and storage plant can fluctuate is 2 percent per minute. And in order to meet that, a different sort of ramp control limit, you'll need different sizes of storage.

So the blue line in this chart just shows how you might expect a PV plant output to suddenly drop away as a cloud passes overhead, and then the orange line would be what would be the maximum ramp that would be allowed, given how you've defined your ramp control limit. And the gap between those two is going to sort of define the size of the storage that you're going to need.

And so just to kind of put some numbers on this, the middle column in here shows the battery duration in minutes that are required to be able to meet different ramp control limits. So if we go to the 10 percent case, that means that this combined PV and storage plant will ramp no more than 10 percent per minute, and the battery size that's required to do that has a duration of about eight minutes. It means that we could fully discharge that battery in eight minutes if we were to run at full output. As we start to go to a more strict requirement, again, if you go down to 2 percent or so, we might require a battery that's more like 45 minutes, or down to 1 percent it's a battery that's longer than one hour in duration. The other thing is that the nameplate capacity of this battery has to be pretty comparable to the solar PV nameplate. These are -- you can expect, again, something -- at an individual plant you can see nearly the entire nameplate capacity drop-off because of clouds. So the size -- the nameplate capacity of that battery is pretty large, and then its duration depends on that ramp control limit.

So here's just an example in the next slide to illustrate what this looks like. So in the red lines we, again, have sort of our PV system that would be uncontrolled, and this is the fluctuations you might expect. On the left-hand side now it's a fairly clear day with not a lot of clouds passing overhead. And then the chart on the right we have a partly cloudy day. And this is data now from Florida, where one of the features is that you have a lot of low, fast moving clouds, so you can see a lot of variability coming out of individual PV plants.

So if we impose this ramp control limit, and in this case I'm illustrating 5 percent, what that means is that the battery is going to act in a way that limits the ramp rate of that aggregate system to the dark blue line that's sort of overlayed over that. So we still have variability, we still have some fluctuations that are happening from the combined plant, but the ramp rate will no longer exceed 5 percent per minute. And in order to do that, that battery is going to have to be discharging and charging fairly rapidly as those clouds are passing overhead, and then that means that the -- and that's what's shown in the middle chart there. And then below that is the sort of remaining energy in the battery so its state of charge, essentially. And you can see that it's moving through various parts of the cycle there, but we're never going to sort of -- we never depleted the battery at all, so even on this fairly cloudy day we move through a lot of the energy part of it, but we don't deplete it.

And so this is a methodology -- this is a fairly simple control algorithm that is something that is implementable today. It doesn't require any advance forecasting or perfect foresight or anything like that. It's just something that you could implement with a battery and PV storage system, and as long as you size that battery sufficiently, you'd be able to meet that ramp control limit that you had specified. Okay. So the next slide.

So our question was to take sort of this control model, a particular ramp limit in a particular PV plant, and combine those things to size our battery and then dispatch that battery so we can come up with what the battery -- what sort of cycling you would have required from that. And then what we did is we passed that off to an NREL tool that's called the System Advisor Model, and recently they've added the ability to dispatch and look at lifetime of batteries in that model.

So we take our minute-by-minute dispatch profile from whatever case we've just run and we feed that over into the SAM Model, which will then look at what that means in terms of the performance of the battery and how -- in particular, how much it's going to degrade that battery as it's run in that way. And one of the things it will specify is that if that battery goes down to 80 percent of its original capacity, we need to replace it. So it has to -- as you cycle it more and more, it's going to degrade the battery, and then we allow that to occur up to 80 percent of its original capacity, and then we swap it out with a new one. And so that means that we're going to have an upfront cost and then some replacement cost over the lifetime.

And we take all that information out of the SAM Model and then come up with sort of an overall cost of these systems, sort of how often they have to be replaced, and then what the costs are per unit of solar.

And slide 24. That's all summarized here. So we have two different cases that we did this with. We have a PV plant that's from the City of Tallahassee. That's in the orange bars that are taller. And then we have a PV plant that's from Jacksonville that's a lot larger of a PV plant. So Tallahassee has a 75 kW plant that's fairly small, and then Jacksonville has this 12.5 MW PV plant. And otherwise we keep everything the same, so it's just that these are two different sized power plants.

And then what the different bars here are representing is if you go from the right to the left, that's increasing the stringency of that ramp rate, so we're trying to make this a smoother and smoother output at the individual PV site. And in order to do that, again, we had to size that battery larger and larger, and having a larger battery means that we're going to end up increasing the cost of that system. And so in order to meet that stricter ramp rate requirement, we will have a higher and higher cost. And the cost that we show here are taking that incremental cost of the battery per unit of solar energy that's put out there.

So you might think about it as if I had a solar power purchase agreement price, you know, maybe something -- these days people are sort of talking about like 30 to $40 per MWh, and that was for an uncontrolled PV system.

If I now say now I need you to meet this ramp control limit, I'm going to specify this ramp control limit, and that person goes out and buys a battery and sticks it on top of that PV plant, their cost will increase by the amount that's shown on the vertical axis. That's our sort of estimate of how much the PPA price would have to go up to make that make sense for them. And then the dots here just show how often those batteries have to be replaced over 25 years.

So I think one of the things that stands out in my mind is that these numbers are oftentimes higher than what the integration costs are that we talked about. So, you know, again, some of the range of those integration costs did go pretty high. They might go, in extreme cases, as high as $20 a MWh, but most of them clustered in that $5 per MWh range or somewhat smaller. And so from what we've seen here with -- you know, it's a pretty simple analysis that we did, so we didn't go into a lot of advanced controls or a lot of different things, but we saw that by doing that at the plant level, it could be pretty expensive relative to some of those integration costs.

And then the difference between the orange and the blue bars is in part kind of, again, comes back to this storyline, is why are integration costs not higher, and that's because of that aggregation effect. And so if we have a small PV plant that's only 75 kW and we're going to chase around every single variation, we're going to do a lot more moving of that battery and we're going to replace that battery a lot more frequently.

On the other hand, if we had a larger PV system, that even within the footprint of that plant we're already starting to smooth out some of that variation, the battery doesn't have to move as much, doesn't have to chase as often, and so it won't be degraded as quickly and you won't have to replace it as often. So we think that the difference between the orange and blue bars are largely driven by some of that aggregation effect even within the individual plant level that can occur. And so if you extend that out and kind of think about what would happen if you go to multiple plants or a balancing authority, you might start to need less and less of a battery to achieve that.

The last couple of slides here -- go to the next slide -- are just to kind of wrap up what we've talked about today. Again, I think a couple of key points just are that we did see that this capacity credit of solar did vary quite a bit, and that often had to do with the differences in load pattern, so sort of understanding what the -- what's driving the load here and how that may differ from what you've seen in the past or from what people experience in different parts of it are important for understanding that. And then for the role of storage, that really understanding how many hours of duration you have really is going to affect the capacity contribution of storage. And then if you start to combine these things within the same power plant, then if they're sharing infrastructure, if they're sharing an inverter and interconnection, that might actually start to limit the capacity credit that you would get if you have batteries that are sized comparable to that PV system.

You know, I showed some evidence that we've seen where by smoothing over a larger footprint, we're able to lessen some of these integration challenges. If that's not an option and you really have to do it at the power plant level, there are some ways to do that. You can sort of do that through specifying these ramp rate limitations, but there are costs associated with that and they may be higher than some of these integration costs. And, again, how strict you want to be on those ramp requirements is going to dictate the cost of that. And that these small batteries are going to be seeing large charge and discharge cycles, and that's really going to cause the degradation of them that is an important factor.

And then the last slide -- go to the next slide, please -- is just to, you know, acknowledge that there are a lot of different directions you can go with this and there's a lot of questions. We're sort of just scratching the surface to get an idea of what are some of the important questions.

And so one of them could be is that there are different ways to control the batteries in these cases, so we're specifying a fairly simple, easy to implement method, and maybe there are ways that you could control that battery that could still achieve some of these same outcomes without degrading the battery as much. And so those might be an upper limit on some of those costs.

And, again, one of the things you want to think about is that rather than trying to control it at an individual plant, that there might be ways to sort of take advantage of geographic diversity to smooth out some of those ramps prior to trying to control that individual plant.

And then, you know, thinking about things about some of the flexibility from PV curtailment is another option that we didn't go into or dispatching from the PV itself, and then looking at what the costs are from those dispatchable generators. That's really where the integration costs come in handy.

And a final thing to acknowledge, I think, is, again, these were kind of different aspects that we looked at, and we never did try to combine or we didn't successfully combine a single battery that's providing both this sort of resource adequacy contribution and ramping at the same time. That's likely something it could possibly do is to get multiple services out of it, and that might then have somewhat different cost indications that what we've shown here. We've sort of done this as independent study so far.

So that does it. And we have a few more minutes for any additional questions.

CHAIR MITCHELL: Question, just to follow up on the last point you made.

DR. MILLS: Yeah.

CHAIR MITCHELL: So it is possible for a battery to serve both use cases?

DR. MILLS: Yeah. I think the -- I think part of it comes from the fact that these are sort of operating over different time scales, that the capacity credit stuff is a lot of what you're going into, and you're thinking about the resources that you need to have on the grid from a planning perspective, whereas the ramp control stuff is really something that you're doing on the operational side, and oftentimes those peak demand needs might sort of align with when you'd be doing this from the battery anyway. So they don't necessarily collide with each other.

And we did a little bit of this trying to dig into that, and we don't have anything to share with that, but so far what we were able to see is that we could achieve pretty much the same capacity credit from this battery if you had like a four-hour battery, plus do that ramping control of it. Now, what the cost implications are and those sort of things, that's what we didn't get into, but it is something that's possible.

CHAIR MITCHELL: Okay. Thank you. Additional questions?

MR. MCDOWELL: I've got one, Andrew.

DR. MILLS: Yeah.

MR. MCDOWELL: I gather from your evaluation of what integration services cost look like compared to utilizing storage to mitigate that cost, there's not really a value proposition there because -- based on your modeling anyway?

DR. MILLS: Yeah. And I think this kind of goes back to what Dr. Taft said, that if you're kind of thinking about only deploying that storage on the edge at that individual sort of isolated system, you really are potentially losing out on some opportunities.

So that is what we found, is that, you know, if you had -- if you're faced with a choice of do I aggregate at the system level first and then deal with whatever is left over or do I smooth everything out at the individual locations first and then sort of, you know, balance the system from there, we see some suggestion that it's cheaper to do that if you're aggregating at the system level.

Now, I think one more aspect to that, though, that we didn't go into is that what if you had storage out at that PV plant, but instead of controlling that storage, just to follow whatever that individual plant is doing and had some way to coordinate across that and say that storage could actually be providing a system asset, I could send it a signal that says, hey, right now it would be really valuable for you to dispatch in this particular way and I could get a service from that battery that would be system dependent rather than, you know, just watching what's the cloud doing overhead right now.

MR. MCDOWELL: Right.

DR. MILLS: And that -- there is a value proposition potentially there, but it does require coordination of some sort of mechanism for doing that.

MR. MCDOWELL: Okay. Thank you.

CHAIR MITCHELL: Any additional questions?

(No response.)

CHAIR MITCHELL: All right. Dr. Taft, Dr. Mills, we very much appreciate your coming to be here with us today, and the information and materials you've shared, it's been very insightful and helpful to us. And with that, hearing nothing further, we will be adjourned. Thank you.

(The proceedings were adjourned.)

____

STATE OF NORTH CAROLINA

COUNTY OF WAKE

C E R T I F I C A T E

I, Linda S. Garrett, Notary Public/Court Reporter, do hereby certify that the foregoing hearing before the North Carolina Utilities Commission in Docket No. E-100, Sub 164, was taken and transcribed under my supervision; and that the foregoing pages constitute a true and accurate transcript of said Hearing.

I do further certify that I am not of counsel for, or in the employment of either of the parties to this action, nor am I interested in the results of this action.

IN WITNESS WHEREOF, I have hereunto subscribed my name this 3rd day of March, 2020.

Linda S. Garrett

Linda S. Garrett
Notary Public No. 19971700150