Event ID: 1905417
Event Started: 2/15/2012 9:50:21 AM ET

"Please stand by for realtime captions."

[Conference to begin shortly.]

Welcome and thank you for standing by. At this time all participants will be in a listen only mode or until the question-and-answer session. To ask a question at that time, Questar then one. This conference is being recorded. If you have any objections you may disconnect at this time. With that is my pleasure to turn the call over to Mr. William Saunders.

I am not Bill Sanders. I Inc. keep Marzullo. I have known Bill for longtime. It's really a pleasure to have him here for the distinguished lecture speaker today. I've known them for longtime, but I'm to read his bio because there's many details here that you should no. We were just talking about thinking of him as in Arizona professor, that was 17 years ago. William Sanders is a Donald Willis professor of engineering and it director of the quarter ended -- at the University of Illinois. He is a fellow -- Tripoli committee on a ballpark continuing. In the working group 10.4 continuing. Has researched includes dependable -- and security. And evaluation on practical infrastructure. The DI be in for structure power grid partake -- initially funded by NSF are he's also a member of the advisory committee. This is just a drop of the things bill is involved in.

Thank you Keith, that's really nice. You didn't tell any of the [Indiscernible - low volume]. Thank you, it's a real honor for you inviting me here, and it's an honor to be here today. I really relish the opportunity to talk --, something I believe passionately in. [Indiscernible - low volume]. I've been thinking about this for a long time, it's something -- security metrics in Washington -- everybody is talking about it, right? I'm going to talk about the beginning of a path. It is a path I think is really important. At path that I think the timing is right to go down. Years ago -- a great problem. I believe it's the right time to go down it today. -- Tells you what we've been doing the last -- years. That is the [Indiscernible - low volume].

Department of Homeland Security -- the last three years. Also had Hewlett-Packard, General Electric and IBM. [Indiscernible - low volume] possible -- security measures. Industry needs them -- whether rightly or wrongly. A combination of working together with academics -- really important. -- Whether we like it or not -- a lot of money on security. It's a big problem. -- As a leadership program in that area for a decade or more. [Indiscernible - low volume]

[Captioner cannot hear speaker.]

This many things we should do to try to achieve that. What our goals are. -- With systems are more secure. A very important question to know -- do not have to have an answer -- absolute quantitative now you -- I think that's very important at this time. And the second question might be how much more secure --?. -- [Indiscernible - low volume]

I think we can make progress, and we are making progress -- security agencies will -- it is useful if it's not perfect. And before I get onto a bit more technical presentation, is when one talks about security metrics, you could mean -- checklist, you can mean kinds of methods that are used in the design and architecture phase. That's going to primarily talk about today. Metrics that are used experimentally in the field, presentation testing, things related to other quantitative testing. All of those are important, all of those are useful. You can make progress in all of them. It's not one or the other, architecture -- no matter how well you implement it --. If it doesn't have architecture flaws -- [Indiscernible - low volume] today I'm really going to focus on --.I can stand over here, but I'm not moving too much anyway. That's what we're to talk about today. So why do we want to use this kind of evaluation? I think I have talked about this in formally already can't to give a good picture view of the security perspective. How there was a relationship between component level and securities and overall system security, that mapping or understanding is complex. We need to look at it. And really if you have a fixed budget, you have choices about what to do. How should you make those security investments? In many cases, when we work with industry, there is a question. Should we invest more in security education, should we invest more and security policies, should we invest more in security technology. And the only way to answer that rationally is through these kinds of assessments. So that is the kind of thing. I'm focusing on the early stage. I'm focusing on architectural design system design I can be fairly technical. Maybe it can be using configuration design, but we still need those other metrics or assessments. I don't know what the culture is year, were the questions need to be at the end or they can be during, at the end. Okay. I need to go forward and keep going. So what is the situation today? We went out and spent a lot of time talking about people in industry. There are two there are two ways of making security decisions. One may -- one way as to why rely on allies individual. And says I think -- these changes need to be made to make it secure. That is perhaps a good way to make decisions, but it relies on that one individual. It's hard to make it scalable. It is hard to audit it later, the know why you made the decision. And it's hard to create more of these wizards that we can use. Another way we see, and I see this in the corporate environment when I work with them, is a more roundtable type discussion. You get people together with different expertise, you get the people who are system administrators and understand about hardware configuration, software configuration. You get the people who understand about training together, the people that understand about corporate policy together. The trouble is, they talk different languages, and again they have different currencies by which they are measuring security. It is difficult to compare alternatives across the board with regard to these things. So these are the challenges copies of the way it's done today. In fact I heard in a corporate setting, we get together in a roundtable, around the glass of a good bottle of red wine, we talk and make some decisions, and we move forward. And so what we really need, if we're going to have let alone the science of security, but the engineering of security, we need ways in which even if they're not perfect, to enable those in 10 years can't to enable those decision-makers can't to make decisions even if they're not perfect. And then later, know why the made those decisions and go back and make changes to those decisions if they need to. That's what we're really looking for. Creating tools to allow these diverse stakeholders to express these different aspects of security. Let them express their concerns, and do it in a way they can work towards a decision about architecture or about design. And do it in a way that's auditable so that six months from now they can decide whether it's this deal the right decision at that day in the future, or whether there's a different decision they should come to about that. Just to give you a feeling of where we are going, and this is a slide you will see again at the end of the talk, but I wanted to show you the kinds of tools I think

we can create. Without going into the details, I will go into the technical details in a minute. This is something called an attack execution graph. It looks sort of like an attack tree, or an attack graph. The boxes are a tax tips. The circles our goals that attackers might be trying to achieve. And the boxes and the triangles are some references -- representation of the state of the system are the lines drawn on their show the out what other quantitative analysis from the approach we have taken, and show how different attackers with different aims, and with different preferences with regard to how they might attack, are most likely to attack the system. Are our numbers behind the Aeros were you see how likely it is these different kind of things happen. I think the numbers in the absolute are suspect, but numbers I think I can argue for you in a relative way, and in a comparative way, are giving us a lot of information about how we need to secure the system or we understand for example, in this particular case, that the arrow that a lot of arrows are going through, I think all to her of protection settings box, that is a place where you need to place your efforts as we try to get the system more secure. We can also understand different kinds of potential attackers will attack the system in different ways. So they're our numbers behind this, it's interesting to look at the numbers and their relative values, but we also get a lot of insight from being able to quantitatively build, blue is the output of the analysis that is done based on an attack execution graph that an architecture has designer built. There's a little bit of related work, and I don't have a lot of time, the thing about security metrics is that there are many methodologies. Some grounded in sort of, if you will, good evidence and some not. I wanted detail about two general classes of methodology that motivated our work. The first is the longtime tradition, starting with the computing committee of fall trees, now attack trees, and attack grass got a try to map out -- with attack grass and attack trees, or an estate-based way with privilege graphs always in which the system, from the point of view of the attacker, is organized with respect making it secure. This is interesting and good work, and I have to say particularly attack trees and attack grass, the dominant way this is done now in practice when you try to do a security evaluation. But they have limitations because they don't explicitly look at how an attacker behaves, and they don't model the attacker. In doing so, they don't take and it account differences, that different attackers might have with regard to ways in which they attack. The second thing they don't do, generally, and privilege graphs do this, but they largely don't have a rich notion of state in the inner relation of the state of the system evolves over time as it's being used, and how that state can affect whether bolder ability is present or not. So you think about the system and its execution, as it's actuating Oryza mission is being carried out, this is a mission oriented system, then there are to be different timing -- timing may be important. We need to represent the timing some Al. The second large class of related work is an adversary modeling itself. Again there is a long and rich history of this. It tends to ignore the system aspects. In a sense, it is based on classical security when you have the first breach the game is over work so they model attackers in a fairly nice way. But, they just ignore the details of how the system is organized making it hard look at how you architect the system. The two that most closely motivated our work is something called mortar or quarter was developed at the national security agency 10 to 20 years ago or it is publicly available papers you can find on that topic. It's pretty interesting work. And the second is called and red, and that is another tool that focuses more on the network aspects. These realize that different attackers will behave in different ways. We try to motivate representation of an attacker behavior at a medium level. Something more detailed -- scripters to a state model that we talked about so much. And something higher level than a vulnerability database which is so fine -- you would say if it's in the database you should fix it, and if you fix it you perhaps don't need to worry about it anymore. That's not are a helpful inter-.

So our work really combines, issues they -- as you see, the state buddies and I will -- state-based model and the attacker modeling that is been used on the other side of the fence, if you will, for quite a while. So what does advise do? The adversary view of security evaluation. It is driven by adversary and now wants us. There is a explicit representation of adversaries. Not exactly would have the series will be, but what their preferences are. Are they noisy, or the quiet, are they worried about how costly attacks are or not. And so forth, the different and the series will have different propensities when you go towards that. It is a state-based analysis. In a sense, what we try to do is combine -- the acceptance in industry and the practitioners feel, of attack grass and attack trees, with a more realistic state-based analysis of privileged grass, but in a way that security practitioners can believe in. They are quantitative. Again, I'll blood explain myself. I don't na&ï;vely think that these are absolute numbers that come out. But these are numbers that died design and configuration decisions are and they have metrics that are fairly high level and relevant to the mission, or relevant to the application or there are a lot of people that talk about what is the security metric. And I would argue that there is no security metric. There are single -- many different metrics are done any for many different purposes. For the purposes of these high level architecture studies, you want metrics that relate to can the system fulfilled the mission, or is the system operation proper with respect to what it was designed to do. As we saw in the trust in cyberspace report, Fred Schneider led many years ago now, does the system do what is is expected to do and nothing else we have a way of quantifying that at a higher level. Just to give an example, I'm not to belabor it. This is a simple, may be small, skater network. Typically in the power grid, SCA D.A., I can't tell you exactly what every letter means, but it's a process control network. It has the sensors in the actuators which refer to the grid. The control network your, hope it is very secure, it's typically behind two firewalls, L. have a connection to the corporate network, there are rules about what goes back and forth their are and then the appropriate network will have a connection to the Internet. If I want to attack this, how might I attack it? Well I might work my way into the Corporation, and I might do a physical attack on the appropriate network. If I can find a way to plug into a jack in the corporate network, and work my way from the corporate network to the SCA D.A. network.. And I might try to go to the corporate VPN or there are two possibilities, Ither to think like an attacker, you might think about doing. We need to represent these types of possibilities in the overall steps that an attack might take in a logical way in order to do the analysis this is a highlight of the overview of the approach. I will walk you through it are showing what we have done in each of these areas. And the in part of the story is will show you a prototype tool that is now used, and has been tested, in about 10 different corporate and government locations. Were still taking a lot of feedback on making a lot of improvements, but we think we're on the right path going forward. The first step in all of this, how we represent the system? How do we represent the system, and that's done through something called a to and attack execution graph. We added the word at execution to a tag Raff to make very explicit that this is a richer notation -- common tutorial notation, that one typically has in that it takes into account the state of the system. It is a dynamic evolution over time. There is an execution in which the attacker works his or her way through the system. And an attack execution graph, what is it? It's a set of as facts that is -- the tax tips. Here's the attacks that this is being corporate access to local physical access. I get to the door, I somehow get on the corporate network. I might come in in a cyber way through the VPN work it takes into account here, Access, skill and knowledge of the attacker. In order to enable or just enable certain steps attackers might take in any point in time. So these three constructs ask as skill and knowledge, can be dynamic. Skill we believe is the least on the timescale of the analysis,

something an attacker has or doesn't have, and it can be black and white, it can be binary, or it can be more continuous value. And knowledge and access are things that you called over time or the attacker may learn something through a step that helps him or her in the next at, and the attacker access really represents place in the system, or the degree to which the attacker is currently infiltrated the system.So the idea of what is new Powell relative to attack grass, is a new idea of having the motion of state, and the skill of knowledge. There are steps that are traditional, they are also richer. And the steps can then lead to changes in the state variable, for example access to the corporate network in this case. And will look at this as an evolution in time or I need to tell you a little bit more about the details now. Here is an attack execution graph and the formal notation really isn't important here, except for it is formal imprecisely which we can prove things about the optimality of the analysis. We can sometimes compute analytical or numerical solutions as well as simulation solution. For now I just want to show you the basics. So what's an attack execution graph, and set up a tax tips for example excess of network using the VPN. And access domains, and that is the represent things like the privilege that I talked about. The set of knowledge items, and again you can define these how you want for a particular system. It is a set of adversary attack skills. Different adversaries will have different levels of skill in different categories. And it's a set of attack goals, and different adversaries will have designers and different proportion to achieve different attack goals are this is what it is, and of course there is a more formal more precise definition that we need to use, and we write the algorithms of how these things execute. Today I'm just in a present this at a high level. An attack step is in a sense the most rich of these constructs. And attack step is really the event if you will in the system that can change knowledge, and access. So we need to talk about how that might occur. What is an attack step? There are some preconditions that say in this excess skill and knowledge, in this current state of the system talk in the attack begin? That as a time that it takes to complete the attack step. And again, I don't know what the time is it maybe a bearable. We can make it a variable, a random variable. But the key here is that we are making our best guess for what these things are. Were doing sensitivity analysis, because we can solve houses of these models at once, and parallel with the system we built to look at sensitivity analysis. And we're documenting what our decisions are, so later we can say we think that was wrong, let's go back and change it or so we have a disciplined, auditable or seizure for how we made these tentative decisions. So this is the time it takes to attempt to step. The cost of attempting to step. The set of outcomes that might come when the step is taken or simply it might just be success and failure. There might be gray areas of the outcomes. Probabilities of those various outcomes. The probability of an attack being detected when a certain outcome occurs. And the next day that represents how the access, skill and knowledge changes, as the outcome occurs. These kinds of things, you will see they are graspable and a GUI-based approach, the document these kind of decisions that can be done in a collaborative way. And then we know the decisions we've made and we can analyze the sensitivity of those decisions to our decisions we make about security design. And we can go back and question them later if we want to. This is kind of a technicality, but every attack execution graph has something called a do-nothing step. This represents the fact that Olympic you point in time to a particular access, skill and knowledge, a state that the absurd may choose is not the right time to attempt an attack step. And do nothing. And I think a good design would force the attacker into the step quite often. So the details are here, but basically, you can set it however you want for the situation but basically a do-nothing attack step doesn't gain you much, as you see there, you want to detected. But you won't make August either. Okay, the next thing to talk about is adversaries. This work is motivated closely by the work that I told

you about done it MSA on mortar. We like philosophically how mortar represented at the series, it did so in a way that didn't represent the system in the rich way that I told you today. Let me dive right in. So what does a necessary look like? And at very scary and adversary start somewhere, just like -- when you do attend test, you may start again from different points, and put the adversary in different places to look at how they can progress. This is akin to that. We have a starting position. The starting excess skill, and knowledge, for that adversary. We have for each attack skill level to the adversary. How good are they at these things? What are the units, units are relative. They are relative to, as you'll see, the preconditions about when an attack can be attempted, can depend on all of these parameters, and also things such as time, the probability of outcome, everything I told you about on the last slide can be state dependent, and dependent on the characteristics of the adversary or an attack goal value function. To what level do they value each of the goals? Might be going after? Different adversaries will value different goals. Attack preference weights. This is very important. This gives the relative importanceto this kind of Avenue. is very , what they place on costs, payoff and detection probability or these aren't values for these things, these are how important is it to this adversary that when the attempt something he gets a payoff. How important is it to the adversary that he is not detected when he attempts something. You can imagine different types of adversaries will have different propensities to these types of things. These are relative weights and you'll see how they come up in the optimization analysis, that can help us understand in different ways different adversaries might behave. I think it is allowed can't repeat the question.

[Indiscernible - low volume]

Great question. So the system does have state, and what I'm telling you. The question is are you considering the system is static, or is the system dynamic cover example moving a moving target defense. What I'm telling you today, there are three aspects of system state, as I said, there is skill, access and knowledge are it is fairly adversary related. The accesses the if you will, part of the system state that the system may be dynamically fighting back, pushing access in the other direction. I'm not that talk about that explicitly in this talk, but these models, the attack execution grass, and the adversary models, can be composed together and in practice where we use this they are composed together with another system model that has its own transitions, where the system observe the access of the adversary, and may take action. I am getting into more technical details than I can in the amount time we have, but the Mobius modeling tool which is implemented as a rich structure from composition, there is a third model, system model that talks about the events that the system may take, and the way in which these models interact is through sharing the view of the state variables. Did that help? So we look at Roe active and reactive defenses in systems in that way. And a moving target would be a perfect thing to look at. Great question. And then there's a planning horizon. The planning horizon says how far ahead does this adversary look. If you knew where I was going, you would see that we're to be doing a game to Redick combined game to Redick's immolation solution approach to solve this. One of the things we need to do is have a finite planning horizon in the game. So you might say oh, standards put this them to get himself out of a mathematical problem about the infinite horizon. Yes, maybe that's what we thought of that in the first place. But as we thought about it more, as we talked particularly with some of the government contacts, the idea of adversaries having limited planning horizons, and the practical observation of that adversary, seems to be bouncing around without too many Eames. They get to a certain state and suddenly they home in on going toward

certain goals. It suggest a behavior that can be indicative, and you can only plan so far ahead, and only -- the game excess skill, and knowledge talk and they start honing in on what their goals are. This has turned out not just to be a mathematical construct, but a very useful construct in saying depending on how far and adversary can look ahead, how effective will they be your we study that extensively in some of the case studies are I am not really talking too much of WC are, but that is an important construct. So that is the adversary representation. The third thing we need to represent is what is the security question, what is the metric we are asking. I won't say much about this, except that it is very general. There is a language for describing these mission oriented goals that can represent anything you want, any function of model state, they can represent thing's about timing of events, how often are certain of attacks debs attempted. How likely is a goal to be reached, what is the relative weight it which goes will be reached, will go be reached before another goal, and so forth. So the rich, reword model basis for describing these kinds of mission oriented metrics are, we show you an example. The last part I want to show you, before I move on to a quick example, is in some sense the core of how this works. We spent a lot of time thinking about, involving this over the last few years. In a sense what we want to do, is Bill the way in which we can composed together attack execution grass, adversary models, and if we want also composed together a system model that looks at how the system reacts as this is going on. By the way we are also studying with NSA support, how to put a fourth model in, that looks at the good human users of the system, and how they impact security. The human user can create or not create vulnerabilities as they use it or there's usually -- really for aspects and I'm only concentrating on two of them today. So how do we actually execute this in a way that we get credible results? Something that someone tell me early on, when I gave a presentation, was I want to be surprised by the model results. One of the criticisms of the early survivability simulation work that was done, was that you had to specify an attacker at a level of detail that you pretty much knew that what the attacker was going to do. If you're not surprised, you haven't learned anything. So we've had to define the level of representation that people were comfortable with operation -- populating these values I just talked about. But they didn't know exactly what the outcome was going to be before they did the analysis. Perhaps they were surprised by the outcome your but then once they saw the outcome, they thought about more carefully and said yes that is credible. We don't have any room for validation that these models represent reality yet, that's going to come over time. But we need models, and I've done this in other domains, in a communing domain in industry. You need models in which they actually give people advice that they didn't know before they started, and yet when they think about that advice it is credible. That it really, they can understand why that make sense. So we're thinking about that both as we defined the level of detail of the modeling, and as we came up with the analysis approach. And here basically is how we did it are at the highest level there is a fairly standard, and I should say there are analytic solutions to these models we are just finalizing on paper now, you can do mathematical solutions with pretty interesting properties because the weather formally defined it at the level Amanda talk about today, this looks like a fairly traditional simulation loop at the higher level -- state. You determine all the available attacks it's in that state are you choose the most attractive of those attacks debs, using to define what I mean by attractive. Used the plastic we select the attack step outcome, and then update the state. The interesting part of this whole the, is how do you choose the most attractive a tax debts? How out of the many things that the attacker could do, the attacker with preconditions is true to be more technical about this, how do you choose which one the attacker will do next? That has to be done in a way that takes into account the preferences, the state, and so forth. So here's the little simulation loop written out a

little more technically. The interesting thing is that this adversary at acquisition. To just talk about this a high level, and I could spend the whole hour talking about the mathematics of how this is done or what we are doing is that each step to win his choices made, each step in that you loop, we are solving a Markov decision process like solution. Were doing and optimization solution that says what is the optimal thing for the attacker to do next, given the characteristics of the attacker. So as it says they are, it says the attractive if the next available attack step is a function of the payoff of the expected states, and steps ahead of him looking in steps forward, and the expected cost and detection of those in steps. So I am looking forward, much like I would in again correct solution, and IM picking or ranking all of the next debs in terms of that attractiveness. Here is the function can't be attractiveness is a weighted combination, without spending time on details, because I want to show the example. But a weighted combination of costs, payoff and detection are and we create if you will, there is a lot of ways in which you make this area efficient, but we in some sense saw the markup decision process optimization problem at each step in the loop of the simulation. And you say all that is very expensive, and I was worried about it being expensive, but it does in a practical way execute very fast. There's a lot of knowledge we can use about the solution at one step, that makes it efficient to solve it again in the next that through her so this is a little animation that shows how this is computed. Basically we have to choose, at this case we have to choose among a set of available next attack step that we can take. We do that by, -- it is simple by just computing the values of the various choices. And we pick the most attractive one. In this case it is the do-nothing attack step. It's more interesting when you have a multilevel attack tree a you have to sort of pushed down and curse up. You make these -- computations. As you are cursing up, in a sense whether his choices, you have to make a choice at a certain level and propagate that choice up, working away all the way up to find out what the next that choice is. If you been paying attention, and you notice the mathematics of there are only certain times in which you can make these intermediate choices. In fact it's a mark of decision process solutions, if you can make the choice at each step. To be technical about this, the solution that we have with this attractiveness function will be optimal with respect to the point of view of the attacker in certain situations. Not in all situations. Basically, it will be optimal if this attractive functions as linear or multicomputer, and it will be totally linear if I had just cost and payoff, or multi-particular if I only had way to detection. What's that saying it doesn't use this attractiveness unction, I can study cases in which I have Austin payoffs that the adversary cares about, but not detection or care about detection and not cost and payoff. There are six, and the sixes basically to represent the detection probability, which turns out to be the non-detection probability in a log way. -- Log way. We look at a geometric and we get around this project -- problem. I thought I'd better show you the example. So what's the practical implication of all this. The implication is we're doing assimilation in a sense, but each step of assimilation, unlike earlier kinds of survivability analysis we are making optimal decisions or worse case decisions with respect to the preferences of the adversaries. There's a case study I want to show you. A case study of the distribution side of the power grid. Is a technology called a pull top we closer. Without going into the technical detail, a pull top of people in the audience know more about this than me. A pull top we closer can be a way to have a second circuit that can be closed got in at this fault in the line that causes an outage on a feeder circuit, the second feeder circuit and feedback first one and help keep the power on to those houses, are those businesses, delivers on that feeder. The question from a security point of view, is is it a good idea. Do we increase the availability of power on the grid, if we put radio we closer is. Do we put aside their feature on these pull top we closer there's and manipulate them from the

control station? The positive advantages are we can avoid outages because we can fix it problems occurred. Potential problem is that we can provide in attack factor for an adversary who might try that control these pull top we closures. We've learned interesting things about how the passwords work on them, how they are part Nader, but that's not the subject of this talk. There are a lot of risk with doing these things as well. Here is the attack execution graph of the system or this is the one without the pull top radio we closures. You can see a mapping between the fiscal relationship and the attack execution graph or this is a realistic model for a typical industrial setting that's been sanitized in a way that we can present it. You can see architectures of this for example in the book by Sweitzer on power grid system architectures. This is how the attackers that we studied, we have studied attackers that had various cost detection and payoff preferences. We studied six different attackers here. Everything from an attacker, they're always foreign governments right, not dismissed it not domestic governments. Attackers that don't have much preference for cost, and maybe equal preference for the detection and payoff. Attackers that might have more preference for cost and other things, for example insider engineer would need to have a pretty big payoff. So we look at how different attackers would behave if they have different preferences with this kind of attack execution graph. What are the metrics? Than average number of attempts that different types of adversaries will make it each kind of attack step or the probability that certain goals will be achieved within a certain period of time and the average time to achieve those goals. You can see how these can be very mission oriented metrics. There is a tool as I said, this is built on top of an extensive modeling tool for system availability and reliability, and performance, funded by the next generation software program, MSF, over many years in the 90s. A tool called Mobius. We built the formulas on top of Mobius so we inherit all of the, if you will, execution power of the Mobius tool. That's widely distributed, 500 or 700 site licenses, used widely now. It's been built upon Mobius. This is how we represent adversaries. XP is a different present them. You see an avid Seri has initial knowledge, initial access, some goals, some relative payoffs that it gives those goals are and then we get results. There are numbers behind these results, but the interesting thing is we find as we study this more more, with more and more systems, the different adversaries 10 to at in quite different ways. It is not like a single adversary will have many different paths all of equal probability. At would be sort of a bad results. Obviously you can construct a system where that would be the case, but for the practical systems we been constructing, we tend to see the different adversaries tended to aid in different ways, that have a lot of cluster and how they behave. We see this result, without the be closer radios, and you see here's the interesting thing if I go back and forth when you add every closer radios, you can see how the attackers really focus on them, and they really, it's an argument for how secure are and how you have to focus your attention on implementation validation of security, on different things depending on what the system. This really helps you focus that next age of security testing, which you done the architectural analysis. There are other kinds of metrics that are interesting. Use the time to achieve attack goals. I want to say right off, don't really worry much about this unit on this access. It is interesting, the different kinds of attackers will have different major goals, and they will take different times to do that in a relative sense. This is all about gaining insight into the system from an architecture and design point of view, as you can see these kind of things. And this graph, will put on quickly, another piece that I didn't put on the slide or interpret it. These are different attack steps. You have to know which attack steps they are for this to be interesting, so I apologize, I wasn't Kohler: did the right way. These are different attack steps, and you can see different types of adversaries will attack different, will try different kinds of attack steps, and

they will try to with different frequencies. That also I find this very interesting information as you try to architect secure systems are what is the future? The question was a perfect setup question, is how does this all fit together. I told you had a model adversaries and how to model the system from an address or a viewpoint. We also need a model both the system itself, and how the system responds. And we need to model humans and how they group humans just like the adversaries to interact with the system or we have a new project started that's modeling cyber human systems. And we're looking at the concept of opportunity of willingness and capability as an organizing construct just like access, skill and knowledge is off from the human side. And we're building formalisms that look somewhere like process models in which show how humans use these systems in our humans can affect their security. So here is an and knowledge the slide. Of course this is the work of many people I many people have done more of the detail work than me. Mike Ford and Elizabeth's Lemay copy as a student. Alyssa that just finished her PhD and has a lecture position at the University now or she's exploring a lot of options. Can Keith is a full-time Senior programmer the builder software tools, and the lead developer of the Mobius tool. We have found in order to have impact of these kind of tools that you can't just -- do the math you have to build real tools people can use. As we said we have thousands of users of Mobius, and many fewer users of the security extension so far that Ken is the full-time keeper of that code. Carol murky from the cyber Defense agency was a collaborator with us on this project. We work closely with people at the NSA, and particularly on the adversaries side of the modeling equation. And we work closely with GE research in the energy division. This is being used to model some of their architectures are as I said, this research was funded by science and technology at DHS, by GE research, by the NSA science and security Center, as we talked over the last years. So what is the punchline? I believe very strongly that the quantitative security evaluation, of this type or of others, there is a whole groundswell of these kind of things starting to come up now. Is useful even it if it isn't perfect or we have to move away from the holy Grail of saying if it's not absolute, if it's not completely accurate, it is not useful or we heard famous statisticians say all models are wrong, so models are useful. This is an act category. We can use these things to make useful the design and configuration decisions that --. I hope we can talk more about this off-line, but better of the decisions being made today, driving us towards an engineering of security. But this is just the beginning of the story cannot be in. In a way to move it forward is to get practical tools in the hands of people that they can try and find out what's wrong with it. We're getting a lot of data. Were to be changing a lot of this in the future to move it towards the practical tool that people can actually use to architect the systems. You need to put those things out in the field, so you can get feedback on how to change it. And in addition to that kind of thing, there's a lot more work to be done, very quick analytics solutions to these models, just what like we have analytics solutions to the fall trees, things we can do their. Just like we have analytics solutions of Markov models, tons of millions of states, how to do that are and a very important thing, all in human behavior side and the system reacts inside. I didn't talk about the system reactions I'd, the capital we pretty much know how to do. Modeling the human behavior side is a much bigger challenge. The human users.So that's what I have today. I don't know how much time I have, but I'd be happy to answer questions. I hope you found that interesting and useful. I will turn it back over to Keith.

I have a question. [Indiscernible - low volume]

In many senses you may not have those probabilities --

So the question was that you may not know many of the parameter values exactly. When you put those in, you may come up to different, at the parameter values change you may come up to different solutions when we are doing the markup decision Rossa solution and ultimately the overall some Malaysian solution. The question is is there a notion of robustness?

There are no formal resolution of robustness that I can make groups of. In analytics of sensitivity to parameters, maybe I can do that maybe I can't. Probably I can do it in some limited cases about I can't do it in general are from a practical point of view of how people are using this now, first alpha users, they are doing very widescale sensitivity analysis. Because of the way all of this is automated, there is a notion of things called experiments and studies, and you can do a study in which you explore parameters input space with many different experiments, in a fairly automatic way. You can get grass that look, for any particular system, I can tell you there's sensitivity is that it's in that -- way, not a beautiful mathematical way. That's where we need to go, and we need to understand that, because even if we don't know them it could be interesting to know how sensitive are those parameters to your decisions because if you don't know the values, if it's really sensitive, you better figure out what the value is. You can use this in reverse also, which is what people been doing in practice. That's a great question.

Sam?

That's why I'm here.

[Indiscernible - low volume]

I have lots of thoughts about this. I have to tell you, that question when we started this work, was paramount on my mind. In fact, if I couldn't get comfortable with in my own mind we weren't even can go forward. So let me tell you what I did very briefly, because is other people when I ask questions. It's a really tough question and I don't have an answer to it yet. What we did, we went around and ask a lot of people, people in government, people who do this on their daily business and government copy Glenn industry. We asked them the question, we validate these models? Had a we be sure that we're doing something useful that doesn't misguide people? And the answer I came to is that really, but label needs to go forward, but the research isn't completely satisfying, people are making these decisions today. They are doing it using a whole variety of methods are and so in a sense, I lower the bar of my expectation of a full validation by saying I believe that I can do better than we are doing today, rather than just I'm going to do it perfectly. And saying we came to come I key word was we are becoming more comfortable with the fact that we're very uncomfortable with the approach. So I just want to admit that first. -- The best answer is I don't know any way in which to validate, if you will, the correctness of the approach of this. Rather than through careful use over time and study. I know that is expensive, I know that's hard to do, but remember we're making these decisions every day. Were deciding whether to make it more secure, spend more money on education to do this and that. What we have to do is we have to try to use these kind of approaches, the document, I worked as consultants in the industry on the power side more than the security side because that was back

in the 90s. You document everything you do and then you go back and look at your decisions. The best I can say, is only over time will we know are in comparison, an usher you*, that in comparison with the other processes, the design code processes, we need to be doing all of these things. In validation, I could've come up today and talk to you about some operational firewall testing things that we are developing into use in the field based on formal method approaches. We need all of those tools. This tool is of the holy Grail of that it is the beginnings of a tool that I think are the directions we should take toward system design architecture process. I'm not to be able to go off and you give me $100,000 and I write a proof and we know it works. It's more complicated than that. That's why we involve industry so closely in this. We get them to try and we get a lot of feedback, and feedback on all sorts of levels. It's a hard question, and I think we have to -- when I came to is it still useful to go forward even though it's a hard question. That's really, we took a lot of time to come to that are there are many questions.

[Indiscernible - low volume]

Obviously you can study that. Yes. So it is a great question. This is a general system analysis question. I worked in a performance filled, the reliability field, the security field always in these kind of system analysis environment. The answer is, this isn't, the tool just doesn't give a magic answer and you do it. It's humans and the insights you gain in the tool, and the humans that make those decisions. The first example, were different kinds of the attackers and different paths, but there are a lot of different paths. What does that tell you? You need to think about which type of attackers are most important to you. I had a mentor, Bill Carter, one of the pioneers and founders of the field of full-time computing back in the late 70s and early 80s when I was a grad student. Bill Carter said to me once, he said it's not about numbers it's about insights. And actually said dammit it's not about numbers, it's about insights are and it's important, were not to take the human out of the equation, but we are going to cut you asked a really good question. I haven't gotten deeply into big security projects where I have helped a team work through that process and security. I have, as a consultant for several companies over a decade in the availability and reliability space, and the way you answer those questions as you use these kind of tools as guides, and ways in which to get data that you then have discussions about. Then you make decisions. That is the way it's going to be here to. That is the way it is in the locker rooms around this place and in the corporate offices to cover those decisions are being made today. At least to my knowledge, it's just they don't have these tools to help them with the decisions. That's I think the best we can do today. A great question.

Time for one more question.

I am sorry, I went on to long and I was trying to be short.

[Indiscernible - low volume]

Great question again. I would guess it was very system dependent as to the number of attackers was important, or the skill of an attacker is more important. You should use this kind of spiritual -- this kind of tool the study that in a particular setting. As to how you would do it using this tool, we thought a lot about this as well. For multiple attackers that are can -- colluding together, and interacting through states and so forth, and have some common view, you model that the even

though I kept using the word attack or in the singular, you can model that type of attack or as a set of attackers, and I just didn't talk about it to keep it simple. We have done those kind of things. If you imagine a system in which there are a lot of attackers that are completely independent from an -- one another, we acting on the system at the same time, if you will by luck or because of some sort of independent similar goals, they are changing the system state in a way that helps each other, or hurts each other. Our current implementation doesn't support those multiple attackers, that is just and implementation detail. -- With three months, or for -- four months of programmer time I can implement that, there's nothing difficult about that in all the theory that we are doing. We just didn't choose to implement it Oeste. -- First. We are going into a next phase which is could be a major new development in the tools were able to release the people, was impending we just got. So we probably will implement that. That is a great question. Understanding again insides, understanding those kind of choices, I don't think we as a community know enough yet. It's good to be system dependent. I guess I am out of time. Thank you all, I had a lot of fun.

[Event Concluded]