

NSF VMware ECDI Program Webinar

CAPTIONERS HELLO CAN YOU HEAR ME Please standby for realtime captions.
TRAN 20. -- [MUSIC] [Music] [MUSIC]

PLEASE CONTINUE TO STAND BY 42587257, THE CONFERENCE WILL BEGIN
MOMENTARILY, THANK YOU. [MUSIC] [Music] [MUSIC]

WELCOME AND THANK YOU FOR STANDING BY, AT THIS TIME ALL PARTICIPANTS ARE
ON LISTEN ONLY MODE, DURING THE QUESTION AND ANSWER SESSIONS PRESS ARE
FOLLOWED BY THE NUMBER ONE, I WOULD LIKE TO BEGIN.

GOOD AFTERNOON, MY NAME IS JIM KUROSE, I AM THE ASSISTANT DIRECTOR AT THE
NATIONAL SCIENCE FOUNDATION FOR COMPUTER AND INFORMATION SCIENCE AND
ENGINEERING. I AM JOINED BY COLLEAGUES FROM VMWARE. FROM VMWARE, DO YOU
JUST WANT TO SAY HI.

DAVID TENNENHOUSE.

CHRIS RAMMING.

SUJATA BANERJEE.

SEAN CROTTY.

GREAT, WELCOME AND WELCOME TO OUR COLLEAGUES FROM VMWARE. AS ALL OF YOU
KNOW, WE ARE HAPPY TO DO THIS WEBINAR FOR EVERYBODY TO LEARN ABOUT THE
JOINT NSF VMWARE PARTNERSHIP. WE GO BY ACRONYMS ALL THE TIME, THIS IS
CALLED ECDI. IT IS PROJECT THAT EXPLORES COMMUTING AND COMMUNICATION
THROUGH A NUMBER OF DIFFERENT PROGRAMS WE HAVE WITH NSF THROUGH COMPUTER
SYSTEMS RESEARCH AND NETWORKING TECHNOLOGY AND SYSTEMS AND SECURE AND
TRUSTWORTHY CYBERSPACE AS WELL AS OTHER PROGRAMS

WE ARE EXCITED ABOUT THIS BECAUSE THIS IS A JOINT COLLABORATION BETWEEN
THE NATIONAL SCIENCE FOUNDATION AND VMWARE. THE GOAL OF THIS IS TO BRING
TOGETHER RESEARCHERS TO FOCUS ON THE QUESTIONS OF WHAT DATA CENTRIC
PROGRAMMING PARADIGMS AND DATASHARING APPROACHES ARE REALLY NEEDED TO
ENABLE COMPELLING APPLICATIONS AND REALIZE THE OPPORTUNITY OF BIG DATA IN
A MOBILE AND IOT TYPE OF ENVIRONMENT.

THAT IS OUR HIGH-LEVEL GOAL AND CONTEXT FOR THIS. THE PROGRAM IS AIMED AT
ADVANCING AND LOOKING AT THE DEPLOYMENT OF SUCCESSFUL EDGE COMPUTING AND
COMMUNICATION ARCHITECTURES BY LOOKING AND EXPLORING DATA CENTRICITY AND
MULTI-TENANCY IN THESE EMERGING INFRASTRUCTURES.

DAVE TENNENHOUSE WILL BE DISCUSSING VMWARE'S OBSERVATION OF SOME
TECHNOLOGY TRENDS THAT LED TO THEIR INTEREST IN THIS PROGRAM. ALSO, HE
WILL OUTLINE ASSOCIATED OPPORTUNITIES AND WHERE HE AND THE TEAM FROM
VMWARE RECEIVE PARTICULAR EXCITEMENT FROM THE ADVANCES WE ARE HOPING THIS
WILL ENABLE.

BEFORE WE GET GOING ON THE TECHNICAL CONTENT I WANT TO SAY SOMETHING
ABOUT PARTNERSHIPS. THIS IS THE SECOND TIME VMWARE AND THE NSF HAVE
PARTNERED ON A PROGRAM. WE FEEL IT IS A GREAT OPPORTUNITY FOR THE AWARDED

PIS BECAUSE OF THIS COLLABORATION WITH VMWARE TO REALLY UNDERSTAND REAL-WORLD CONTEXT AND SOME PROBLEMS THAT ARE FACED IN THE WILD, IF YOU WILL.

IN THE INDUSTRY, THE RESEARCH IS PRE-COMPETITIVE, BUT THE OPPORTUNITIES FOR DIRECT INTERACTION WITH VMWARE PROGRAM DIRECTORS ON THESE JOINTLY FUNDED PROJECTS WILL PROVIDE A VENUE FOR ALL OF US AND FUNDED PIS TO UNDERSTAND AND ENABLE [TECH] TRANSFER. PERSONALLY, AS A PI AT THE UNIVERSITY OF MASSACHUSETTS WE ARE FUNDED BY OTHER PLACES AND IT HAS ALWAYS BEEN REALLY GREAT TO HAVE THESE INTERACTIONS WITH INDUSTRY.

It gives us a view into something we in the academic community cannot see on our own, these partnerships are incredibly important. It will really give the PI community a view and perspective we can't get on our own. Maybe I will put in a plug here for NSF and the foundational work we do and I will hand things over to Dave in just a second.

As a faculty member, if you ask me what NSF does I would say "they do research". If you look at what we call the organic act which founded the National Science Foundation, the goal is to fund basic research but also quote, to promote the progress of science. Also, to advance the national health and welfare, secure the national defense and for other purposes.

Our mission is about fundamental research and the contribution that it makes to economic competitiveness and national security. One way to achieve this is through technology and working closely with research folks from industry.

If you get a chance over the weekend and you haven't read a report called Restoring the Foundation [Resource Center], it's really a great report. It came out about three years ago and calls for new public/private partnerships. A collaboration between industry, academia, and the federal government. We think this collaboration we launch today with VMware and our past ones with VMware and other industries and with the PI community have been great examples of that. We are really looking forward to this and continuing strong collaboration between NSF and VMware.

All of us here at NSF are excited to kick off the second joint research effort with this meeting to discuss what this will be about. Let me welcome people again. I want to give a shout out to folks at VMware. It has been absolutely great working with them. I will turn things over to Dave Tennenhouse, the chief research officer at VMware.

(TENNENHOUSE:>> Thank you, Jim I was happy for you to put that plug in for NSF. The Foundation is at the forefront of computer science research and it is the steward of CS research, the long-term shepherd of research in the U.S. It is a real honor for VMware to partner with NSF and the research community. In this effort to significantly advance our approach to computing.

You were good enough to thank the VMware folks, in addition to thanking them for their help with the solicitation I would like to express my appreciation to the unparalleled team you have at NSF. Many of whom I have known and worked with over the years dating back

to my time as a researcher, including, Erwin Gianchandani, Ken Calvert, Darleen Fisher, Mimi McClure, Samee Khan, Jack Brassil, and Meghan Houghton. and many others I haven't mentioned.

Although edge computing is not an entirely new field, it is poised to take on significantly increased importance as we build out the Internet of things and bridge gaps between the physical and digital worlds. This is a big deal. There are many untapped research opportunities; a central challenge is to lower the barriers to the rapid development and deployment of innovative edge applications.

Given the recent advances in computer networking infrastructures and the growing degree to which they are software defined, it is a time to look at the infrastructure that underpin the edge applications. I would like to briefly provide a VMware perspective on the context for edge computing. We have seen many of the same research challenges Jim identified.

First is exponential growth in the number of network devices and users which causes an explosion in the amount of data generated at the network edge. It is predicted by 2020 that 75% of enterprise generated data will be created and processed outside the traditional data center or cloud. The ability to secure these devices and not data and to do so at IOT scales is not an option, we need to get there.

There is considerable hardware and software at the edge, devices, datatypes, and application services and other components. Some of that is inherent as devices and their data must be specialized to specific applications. The bulk of that is a historical artifact of the numerous vertically filed markets in which IOT has evolved.

Over time we anticipate that reusable infrastructure and APIs come to the edge and replace many hardware components. This will dramatically reduce cost, increase productivity, and make edge computing more reliable and secure to manage. Edge data often pertains to individuals and can be used to infer sensitive information about them. Thus, frameworks for security, privacy and regulatory compliance will be important if we want to see large-scale adoptions of the applications at the edge.

>> If edge computing is to succeed we have to make application development significantly easier, even though requirements of the edge will be more demanding. For example, there will be new types of interactive edge applications that will be latency sensitive and require fast access and real-time analysis of the generated data. We need to make it much easier for researchers and developers to quickly build and experiment with new edge applications.

There are vertical application and infrastructure silos that currently exist and developers often have to create the entire application pipeline from sensor deployment to data generation to analytics to control functions. Everybody keeps having to do the same thing over and over again. This is not only wasteful of data collections, storage,

transmission and processing resources, but also raises the entry for new application development.

A promising approach to breaking these barriers is the data centric architectural approach--the core of the solicitation. This will allow us to remove unnecessary heterogeneity while allowing the analysis to be tailored. This research may open the door to simple yet effective abstractions to build applications. It provides the right abstractions for data, programming, and infrastructure virtualization across the distributed edge cloud infrastructure that can be shared by multiple stakeholders at the same time and can be a game changer.

We believe we are at an inflection point where data will be abundant and unlocking its value through resource and data sharing while maintaining security and privacy will be the key to enabling a large number of novel applications and services.

I want to comment on how this partnership with NSF aligns with the values of VMware. Although we are an enterprise software company today we trace back to government supported research. We have a special obligation to give back to the community and look forward to partnering with NSF to help create innovative edge infrastructure and computer applications for the nation and the world. Thank you for joining us on this journey. I would like to turn it over to Ken Calvert, the director for the Computer and Network Systems Division in CISE.

KEN CALVERT: I want to echo what Jim said about what a pleasure it is to work with the group from VMware and say I am excited about this program. Let me recap the drivers for edge computing. We want to put resources located close to the endpoint to reduce latency and to deal with high endpoint density and compute and storage from impoverished devices.

It requires time-responsive allocation and real-time management of communication, computation, and storage resources. At the same time, many applications that have significant potential for industrial and societal impact bring a number of challenges including tension between privacy preservation and data sharing across applications.

Let me talk about the architectural focus. Today a lot of applications, especially ones that belong to what we commonly call the Internet of things, consist of a device at the Edge connected wired or wirelessly and communicating with an application program running in a data center and the cloud. The device might be a phone, car, various appliances, etc.

This approach enables us to provide security through the TLS and SSL protocols which protect the channel, but it results in the applications being siloed. The data from a device typically goes to one application, this makes it difficult to share data across applications.

The Edge computing phenomenon can be described as a three-tier architecture. We interpose the edge cloud, or clouds between devices and the centralized cloud data center. The Edge infrastructure which resides in proximity to devices, may consist of an IOT Gateway, a content distribution center, a cloud -- cloudlets, or something like a micro data

center. The organization of the edge computing infrastructure needs to support the provisioning of compute, storage and communication resources near the edge to enable applications with special requirements.

What are the important characteristics of this set up? The backend cloud data center, the applications running there, the edge cloud, the applications running in the edge cloud and device may be controlled by different entities and they may have different interests. We refer to this as the multi-stakeholder characteristic.

If you look at the bottom part of this figure [SLIDE 6] there is the dark gray cloud provider, the gold, the red application, the blue application, the purple program running in the Edge Cloud and you have two different devices. You have a potential total of seven or more stakeholders. The second important characteristic is [that] the backend cloud, the Edge Cloud, and device will be running multiple applications. Those may be provided by different entities and have to share resources. This is what we call the multitenant aspects. -- We are hoping to get an architecture supporting ecosystem that gets away from the siloed paradigm we see today and it enables multiple stakeholders and tenants to share multiple types of data across multiple applications in a controlled scalable way.

Of course, privacy and confidentiality must be respected. In a sense, we look at data as the narrow waist to the architecture, to enable dealing with this complex environment. Let me pause here and note that some of the NSF future internet architecture projects have taken an information centric or data centric approach. One may even use the phrase narrow waist in the same way we just did. Building on those projects is one possible approach to this problem.

I want to stress that is not the only one and we are also interested in approaches that are backward-compatible with today's architecture and can be layered atop existing work -- network or other services. The challenge is to enable the great potential of the edge computing paradigm by developing components of an architecture.

One example might be to, say, dynamically or adaptively control complex domains like a smart building. That problem might be approached through a programming paradigm or language, a runtime environment, some kind of data sharing framework, data manipulation framework and so on.

In summary, how can we enable data sharing across devices and applications in the context of this three-tiered architecture while preserving privacy and considering the interests of the various stakeholders? I will hand over to Dr. Samee Khan for more information about potential research directions.

Thank you, Dr. Calvert.

In this solicitation, we request the proposers that they must take into account the multi-stakeholder context of futuristic edge computing environments. More specifically, environments that enable service and device interaction, environments that control exchange of data,

environments that support extensible architectures for a rich, multi-tenant integrated hardware and software ecosystems. In this solicitation, we expect the proposers to address the need for security and privacy as first order design and operational goals. In this solicitation, we have identified three research vectors, namely: system architecture, programming paradigms, and security, privacy, and data sharing, which we believe offer several potential key areas of work. However, the three research vectors are absolutely not meant to be prescriptive or to limit the nature and scope of proposed approaches. We want to make clear that proposers should feel free to propose solutions that address broader edge computing challenges in novel and innovative ways that may not be captured through the three research vectors, which we will describe shortly.

In the system architecture research vector, we are encouraging the research community to envision edge computing architectures that can enable elastic, dynamic, and responsive mechanisms for data and computation movement as needed and as appropriate, and possibly at runtime. To envision such an architecture, one must consider system components and interactions between them, consider issues in operative system for the edge infrastructure, consider effectively and efficiently utilizing a heterogeneous set of hardware and software stack, and consider issues in system abstraction.

In the programming paradigms research vector, we are encouraging the research community to envision programming paradigms that can address multi-tenancy, service interaction, and infrastructure sharing across multiple stakeholders. To envision such paradigms, one must consider developing the right set of programming abstractions to assist in efficient and effective management of the three-tiered edge computing infrastructure, one must consider developing the right data curation techniques to handle all types of data, such as streaming data, time-series data, filtered, and unfiltered data, and one must consider developing paradigms that can in an elastic fashion, provision for the division of computation.

In the security, privacy, and data sharing research vector, we are encouraging the research community to envision solutions that can ensure edge computing infrastructure data protection, irrespective of location, irrespective of data ownership, and irrespective of data types. To envision such solutions, one must consider revisiting data isolation mechanisms, one must revisit the notions of security and privacy given heterogeneous data types, where the meaning of security and privacy may change from one data type to another data type, and one must consider embedding security and privacy into the system architecture and programming paradigms.

In this solicitation, we are encouraging the research community to anchor their proposed work in one or more application domains of societal relevance by considering the intersections of multitude of industries and service providers with a focus on data and multi-tenancy, as emphasized earlier, when we were describing the three-tiered architecture.

In this solicitation, we have listed a limited set of illustrative examples of potential application domains. Each application details an environment having a rich multi-modal data with multiple stakeholders and a variety of real and non-real-time system management and control. We encourage the proposers to have a careful look at the limited set of illustrative examples of potential application domains, as provided in the solicitation. However, the proposers are encouraged to develop novel applications or services (other than the limited set of illustrative examples) that uniquely benefit from data-centric three-tiered edge computing architectures.

In this solicitation, we expect the proposers to develop prototypes of their proposed approaches to explore implementation aspects of their designs and to empirically demonstrate the effectiveness of their solutions.

We want to emphasize that we do not want the community to reinvent the wheel; where possible, and to avoid unnecessary re-invention, proposers should leverage existing software tools and frameworks where possible. Some example of such frameworks (and this is not an exhaustive list) are VM-based cloudlets, EdgeX Foundry from the Linux Foundation, and Open Edge Computing. However, if the proposers find that the existing frameworks have fundamental shortcomings that cannot be overcome by a build, then the proposers are welcome to develop new frameworks. But the proposals must explain the gap and justify the need for developing something entirely new.

As emphasized earlier, this solicitation seeks to advance the deployment of successful three-tiered architectures by exploring data-centricity and multi-tenancy in emerging edge computing infrastructures. The project research should advance theories, algorithms, architectures, prototypes of system components, prototypes of applications, testbeds, and evaluations in the context of emerging edge computing infrastructures. We want to make it very clear that as a deliverable, production-level code is neither expected nor encouraged, nor is the complete end-to-end integration of all research results. However, clear planned synergies between sub-efforts is highly encouraged and highly desirable.

We also strongly encourage the proposers to be precise about the research objectives, research plans, and plans to achieve the proposed broader impacts. More specifically, the proposers are encouraged to leverage on the research results, such as open source releases as a means of building community and experimentation, or testbeds that encourage participation in relevant domains to achieve broader impacts.

At this time, I would like to introduce Dr. Darleen Fisher to continue with the presentation.

For the rest of the presentation we will go over solicitation requirements, the review process and the management of the awarded

projects. We will have an open question and answer session at the end. You can find the solicitation, NSF 18-540, by a web search and it is also posted in the CISE funding opportunities webpage. You should read it carefully as requirements may be different from other solicitations. The proposal deadline is 5 PM local time on May 22 of this year. Please submit them early to your sponsored research office with plenty of time to catch mistakes and make sure you have submitted the most recent version. NSF and VMware expect to award approximately two projects and each one may request up to \$3 million over three years. Awards will be made by late summer of 2018.

Who can submit? Institutes of higher education [IHCs], this includes universities, two- and four-year schools as well as community colleges. I direct you to the special instructions for international branch campuses of IHC's and in particular, they may submit only if the proposed activities cannot be performed at the U.S. campus.

Individuals may participate as PI, co-PIs or senior personnel on no more than one proposal in response to the solicitation. If any individual is found to be on two or more proposals, only the one with the earliest submission date will be accepted. Make sure every individual on your team has consented to participate on the project and is not listed on any other proposal.

We look forward to manageable-sized teams including individuals with expertise that is needed to successfully conduct proposed work. For NSF research grants, we expect to see graduate students as part of the project. You may include software engineers and/or programmers on the proposal as needed. You may have up to 20 pages for the project description and we will check for font, spacing, and margin limitation compliance.

Please include a Gantt chart with major tasks, milestones, and interdependencies. Don't forget the required Postdoc Mentoring Plan if you are funding a Postdoc. You may have up to two pages for the collaboration plan; you must include it in the proposal. The collaboration plan is an important document that notes the appropriateness of each member of the team, including their areas of expertise. It describes the overall management coordination mechanisms and interdependencies of tasks. This helps to clarify how the tasks hold together and the whole project comes together.

Finally, there should be a reference to budget items that show support for collaborations. We expect you to include funds to make it real. Please read the solicitation for more details. You may not in NSF proposals include letters of general support for the project or for its expected outcomes. You may include letters of collaboration pledging, for example, data or access to research facilities.

There is an intellectual property statement that can be found in the solicitation. You can look at the details, but generally the program will operate under a public dedication intellectual property model. The solicitation specifies which open source licenses meet the program objectives, and state the requirement that research results be publicly

disseminated through an access complaint repository in accordance with NSF's public access policy.

Let me emphasize we will follow the NSF process of merit review for proposals submitted to this solicitation. The proposal will be reviewed by a panel and NSF program directors may request additional reviews for a project as appropriate. The proposal is reviewed as we always do according to the intellectual merit and broader impact review criteria. There are additional solicitation-specific review criteria found by reading the solicitation.

NSF will conduct the panels, VMware team members will attend but they will be observers. They will not be involved in conducting panel reviews. NSF and VMware will need to discuss proposals and decide whether to have reverse site visits, which will be virtual by WebEx or some other technology, for the top-rated proposals.

[We may decide to have reverse site visits for top-rated proposals.] At a reverse site visit, individual teams are called upon to answer concerns raised by reviews, NSF program directors or VMware partners. Then we [NSF and VMware] will [jointly] make final decisions.

The solicitation specific review criteria are the following:

- the extent to which the projects have a system perspective --a coherent whole that is more than the sum of its parts. We don't want to see individual research topics that are "stapled together".
- includes an evaluation through demonstrations and a prototype of the components and at the systems level.
- The extent to which projects have well integrated teams of researchers, and
- demonstrate concrete plans to impact the broader industry and
- use existing components and infrastructure where appropriate.

For the funding model the projects will be funded directly by NSF and VMware through separate funding instruments. After the awards are made, the project may receive supplements either from NSF or VMware. Please note, NSF expects the entirety of all awards to be used to support the proposal as negotiated in the awarding proposals.

In terms of program management, VMware and NSF will jointly manage projects while following guidelines governing each party. NSF and VMware will designate a program director for each partnership project award. The program directors will jointly oversee the execution of the projects. VMware will work with NSF as active participants such as advising researchers on technical issues.

VMware will conduct annual on-site reviews and may invite awardees to site visits at VMware, accept invitations to attend site visits at the awardees' institutions, and potentially host student interns. VMware may separately fund its own personnel to directly support the partnership research, part-time or full-time, with the institutions' awarded projects.

In summary, [slide 25] shows how to create a title for these projects so we know the proposals are for this program and can capture them. On this slide I have listed the required supplementary documents. A collaboration plan is typical for a large project. Please read the solicitation carefully and also look at NSF PAPPG for 2018 or NSF 1-18 for overall guidance on writing proposals.

Full proposals are due May 22, 5 PM submitters time. Also, I have listed a "save the day" date for site visits should you be one of the top-rated proposals. We may have reverse site visits and we have reserved July 13 as the date for those. It will be a virtual meeting.

We are now open for questions.

Thank you, at this time we will begin the Q&A session, if you have a question, press star followed by the number one, recorded in clearly and I can introduce you as well as one question. You can go back into the queue and press star one again if you have another question. Again, if you have a question please press star followed by the number one.

I am showing no questions at this time, is another reminder if you have questions from the phone line please press star followed by the number one. One moment for the first question.

Q: >> Hi, thanks for the great introduction. Have a question on how to balance the [theory and practical systems research]- are their preferences on the proposal emphasizing one over the other or any general description?

A:>> I think we want to see prototypes and would like to see things that are clean, but not necessarily too high a level of abstraction. Don't obstruct away all the hard parts, that would be my answer. Darleen, do you have anything to add?

A; I think we are looking for a systems view, to have pure theory would not be appropriate. I think this is more on the practical side than on the theory side.

A: That is a good question. To strike a balance, I think we will leave it up to you on how you want to go ahead with convicting -- convincing the panel which way to go about that. If you read the solicitation, it is clearly indicated it has to be some level of a practical system. Keeping that in mind, you can strike a balance between theory and practice.

Q: Thank you and show the possible application be focused on a specific application, or is it up to us to choose edge computing applications?

A: Again, as I mentioned in the solicitation and presentation, we have indicated a few applications that we think are very important to consider. However, if you think of another application that is of critical need or has a broad societal need, please go ahead and consider that well presenting research to the panel. You don't have to address all of those applications, you need to find the sweet spot. We are not looking forward -- for an architecture tailored to one specific

application, will -- we need to have multiple applications in this contact context.

We expect you have some number of examples to show how your approach supports some application in reasonable detail. Maybe you don't go into all of them at the same level of detail, are you go into a couple at a deeper level. >> The examples were presented as ones which were very rich, they are not like a single application. They are more like an application environment, where the solutions would make sense against a rich set of - the word application. And trying to think of another word.

Our next question is from [NAME REDACTED] >> I had a couple questions. One is, there are a variety of ways to choose tenancy -- given VMware's involvement and their expectation -

[NSF:] we can't hear you. I'm sorry.

Can you hear me now?

Better.

Q: It involves multi-tenancy, is there a variety one can achieve that? I am wondering because of VMware's involvement; if the expectation is that virtualization will be a solution we bring to the table.

>> This is Chris. I think there are a number of ways you can achieve the goals of virtualization such as isolation and the encapsulation of various properties. I don't think there is any presumption about which approach to achieving these kinds of goals you should -

Thanks, Chris.

[NAME REDACTED] had a couple of questions.

Q: Just this one more. When I think about this system embedded in the world, one resource I think about is energy and power. It wasn't mentioned in the presentation of solicitation, is that considered a resource that we may pay attention to?

A: Absolutely yes you can. If we haven't mentioned it, I think it was presumed.

Okay, thank you.

>> Karen, next question your line is open.

Hello, can anybody hear me?

Yes.

Q: I have a couple of questions. The first is since VMware is involved in the selection process, can I interpret it as the award will be selected largely based on VMware's preference? And the second question is the

project will be funded as a regular NSF research grant, will VMware claim any ownership to the intellectual properties?

A: NSF will use the normal process for review and that is stated in the solicitation. After we go through our process and VMware will conduct its own review internally, then we come together and come to consensus on what will be funded. I think the answer to your question is no.

The second question was about IP, so NSF will use its standard intellectual property policies.

Chris, do you want to address the VMware IP approach?

A: Actually, this was a public dedication solicitation. The idea here is we are looking for our research to be implemented and made available under open source licenses and there is a preference for the results to be available to everybody. There is no intent to provide this research to be restricted in any way.

A: I wanted to comment on the first part of the question. This is our second solicitation and experience doing this. We were looking at basic research and the best basic research. VMware looked to us as if they were put on NSF [basic researcher] hats. If they put on VMware hats, I didn't see it. It is very important we are looking at really outstanding research projects.

Q: Thank you. I have one more question. It might be related to the previous question somebody else asked. So the definition of - I am wondering if we proposed to develop a virtualized environment based upon tentative computing solution or different computing paradigm, would that be of interest to VMware?

A: If it sits within the scope as outlined in the solicitation I think it is of interest for this program.

It doesn't have to be virtualization based upon like CPU or an existing, well-established computing infrastructure, as long as it can benefit as computing and Edge infrastructure in general that would be within the scope?

A: Would you ask that question if it was just NSF putting a solicitation out? I think that is the answer.

A: I want to go back to one of the questions asked previously. I think the goals here are to identify the properties you are trying to support with your ideas and venues whatever technology or approach that is the right one. I think there should not be any presumption on that specific kind of approach.

Thank you.

Our next question is from [NAME REDACTED].

A: It caught my attention, the part about making an impact on the broader industry. This is something I often find is an issue, the idea of

engaging with forums such as IETF or other industry or academic forums. There are two problems with that, often that works -- work really starts after the research grant, after doing prototypes, and publishing validation. >> Also it's not clear whether that is considered funded within the scope of the project and it's sometimes difficult to figure out where to get funding for those kinds of a studies. If we are expected to do something other than scholarly publication and open source software and documentation, then I am interested in knowing, is that okay to be part of the budget?

A: I don't think it is expected that you will get to an Internet draft stage by the end of the project. I think one of the advantages of working with a company such as VMware is they have a broad view into the industry and maybe they can help you figure out how to navigate that question.

Q: >> As part of the proposal, is it then - I am trying to figure out what we would say in the proposal since we are developing this without collaboration with VMware. >>

A: The standard VMware -- NSF criteria concerning broader impact has many components. One of which, could be industry impact. The proposal has specific ideas that strengthen the understanding and respect to that criteria. Those activities you propose would be an integral part of the proposal. Again, I don't think there is a prescriptive intent to do things that could not be supported under the funding you provided. >> The work can be evaluated in this to mean intellectual and broader impact criteria. It is a trade-off in terms of how you design your proposal.

>> That helps, thanks.

At this time I am seeing no further questions.

I am seeing no further questions.

Going once, twice, thank you all for your participation. We will report to seeing the proposal bill coming into us.

Thanks everybody.

Thank you, this does conclude today's conference call, he made his connect -- you may disconnect your lines and thank you for your participation. >> [Event Concluded]