

WEBVTT

1

00:00:33.870 --> 00:00:38.040

Matt Mutka: And start this off, I would like to turn this over to

2

00:00:39.570 --> 00:00:46.920

Matt Mutka: Leading leader of our size directorate, the assistant director Mark grid or Canarsie Margaret, please.

3

00:00:48.360 --> 00:00:54.750

Margaret Martonosi: Thanks very much, good morning and good afternoon to everyone. I'm so glad you were able to make it here today.

4

00:00:55.320 --> 00:01:04.890

Margaret Martonosi: My name is Margaret Martin know see as was mentioned, I'm the Assistant Director I lead the Computer and Information Science and Engineering Directorate at NSF

5

00:01:05.400 --> 00:01:13.260

Margaret Martonosi: And I do so as a rotator from Princeton University where I'm currently on leaves, but where I've been on the faculty for six years.

6

00:01:14.520 --> 00:01:28.950

Margaret Martonosi: In those 26 years of research I've had the wonderful opportunity to partner with industry and many, many points. And so it's thrilling to get to be a part of this webinar.

7

00:01:29.580 --> 00:01:39.180

Margaret Martonosi: Which is a very strong partnership with industry and NSF. So right now, I believe that on the zoom. We have

8

00:01:39.930 --> 00:01:48.540

Margaret Martonosi: Matt Mocha who you heard briefly Darlene Fisher Alex Jones and Vaughn layman Eric room fund and go deep saying from NSF

9

00:01:49.110 --> 00:02:01.950

Margaret Martonosi: And we're very pleased to also have from VMware David Tennant house Chris ramming Victor fewer you and I think perhaps some others joining us. And I want to give a big welcome to my VMware colleagues.

10

00:02:03.540 --> 00:02:17.970

Margaret Martonosi: Thanks also and welcome to all of you out there in the webinar for joining us here to learn about this joint NSF VMware partnership on the next generation of sustainable digital infrastructure and g SDI.

11

00:02:18.600 --> 00:02:32.640

Margaret Martonosi: So this new NSF VMware solicitation that we're discussing today and that will feel questions on is intended to bring together research in these and other programs to focus on one big question, which is

12

00:02:33.450 --> 00:02:43.950

Margaret Martonosi: What fundamental and systematic approaches in measurement and design and development and management of digital infrastructure resources and workloads.

13

00:02:44.490 --> 00:02:53.580

Margaret Martonosi: Will enable significant progress towards maximizing sustainability of the digital infrastructure with minimal impact on

14

00:02:54.120 --> 00:03:02.280

Margaret Martonosi: The things that we normally worry about including programmer productivity and the infrastructures performance and scalability.

15

00:03:02.760 --> 00:03:07.770

Margaret Martonosi: So the idea of this program is to foster transformative research and those fundamental

16

00:03:08.250 --> 00:03:19.080

Margaret Martonosi: And systemic approaches to bring dramatic increases in the environmental sustainability of the digital infrastructure and to lead to practical methodologies and tools.

17

00:03:19.740 --> 00:03:35.580

Margaret Martonosi: So where are we going to go next. David Tennant house is going to discuss VM. Where's observation some technology trends that led to their interest in this joint program and the associated opportunities and excitement for the advances that we hope it holds and enables

18

00:03:37.620 --> 00:03:47.010

Margaret Martonosi: What I can say, first and foremost is the degree to which this kind of partnership provides for important interactions between NSF VMware

19

00:03:47.370 --> 00:03:55.710

Margaret Martonosi: And the order p is to understand real world context and problems that are faced by industry, even though the research is pre competitive.

20

00:03:56.160 --> 00:04:06.060

Margaret Martonosi: direct interaction with the VMware program directors on on these jointly funded projects can provide a venue for understanding and enabling tech transfer

21

00:04:07.080 --> 00:04:21.540

Margaret Martonosi: As a researcher, as I mentioned, these kinds of interactions are something that I found valuable in the past because the academic research community off on its own. Can't pose an answer all the questions. It helps if

22

00:04:23.340 --> 00:04:34.260

Margaret Martonosi: We can do so in partnership as you've seen with our new director, the notion of partnerships is deeply important to NSF, as well as to our researchers to industry into our nation.

23

00:04:34.770 --> 00:04:49.320

Margaret Martonosi: So NSF was founded to do basic research, but also to promote the progress of science to advance national health prosperity and welfare to secure national defense and so forth to do all of that we need to actually transfer

24

00:04:50.310 --> 00:05:02.880

Margaret Martonosi: Work the full spectrum from basic research to applied outcomes. And so there's plenty of people that are talking about the role of these kinds of partnerships and we're very excited about them. Next slide please.

25

00:05:05.370 --> 00:05:16.050

Margaret Martonosi: So through this solicitation. We're hoping to foster this multi disciplinary research community at the confluence of networking and computer systems.

26

00:05:16.860 --> 00:05:26.700

Margaret Martonosi: We also hope, as you heard me say, and has you'll hear in the rest of this webinar to transition research results from funded projects into engineering practice to test them.

27

00:05:27.720 --> 00:05:38.670

Margaret Martonosi: In reality, so the partnership begins begins the partnership continues or wonderful collaboration between NSF and VMware, this is just the next chapter of it.

28

00:05:39.300 --> 00:05:51.270

Margaret Martonosi: To combine NSF experience of developing and managing successful large and diverse portfolios with VMware strong capability in the area of sustainability. So I always like to

29

00:05:51.690 --> 00:05:57.330

Margaret Martonosi: Be in the position of kicking off these meetings and then sticking around to hear the rest of it because it's so exciting to me.

30

00:05:57.630 --> 00:06:15.090

Margaret Martonosi: I'm very excited to kick off this meeting that marks the third joint research effort by the VMware and NSF partnerships. And with that, I'm happy to turn it over to David Tennant house who will have the next word he's chief chief research officer at VMware so David. Thanks.

31

00:06:19.110 --> 00:06:35.370

David Tennenhouse: Thank you, Margaret, and it's an honor and really exciting to to be here. I've actually been just really keen to see us launch sustainability effort with NSF, the National Science Foundation has been at the forefront of basic and applied research and

32

00:06:36.660 --> 00:06:46.290

David Tennenhouse: We're just thrilled to partner with you and the whole size research community in this effort to really advance our collective approach to sustainable computing

33

00:06:47.790 --> 00:06:56.550

David Tennenhouse: In addition to thanking all the VMware folks who've been helping with the solicitation. I want to express my appreciation to the great NSF team.

34

00:06:57.420 --> 00:07:07.920

David Tennenhouse: We've worked with many of you over extended period of time around, Jun Chan Donnie, you gotta go deep saying Eric are involved, Darlene Fisher Matt mode.

35

00:07:08.670 --> 00:07:17.730

David Tennenhouse: And one layman Alex Jones and kitchen patrol var you know some of the people that have been involved and we're really excited and enjoy working with them.

36

00:07:18.780 --> 00:07:37.680

David Tennenhouse: As you mentioned, this is actually the third partnership between NSF and VMware and you know the earlier ones began in 2016 2018. So we're kind of on a nice cadence here. So I'd like to sort of just say a little bit about VMware is kind of history in in sustainable computing

37

00:07:39.030 --> 00:07:48.120

David Tennenhouse: You know, first and foremost, our founders had a sort of strong personal bike to work type of interest in sustainability. So it was kind of built into our DNA.

38

00:07:48.870 --> 00:07:56.670

David Tennenhouse: Early on, and then as things turned out virtualization turned out to be a great tool for sustainability, because

39

00:07:57.600 --> 00:08:05.520

David Tennenhouse: Virtualization drove server consolidation allowing people to consolidate many workloads onto one server so that now saves

40

00:08:06.270 --> 00:08:14.070

David Tennenhouse: All of the energy that would go into running that server cooling the servers, you know, all the servers that you don't have to use

41

00:08:14.760 --> 00:08:25.440

David Tennenhouse: When you consolidate those workloads is saved and all the embedded carbon all the energy and carbon that would have been expended to create those servers to actually build the hardware.

42

00:08:25.800 --> 00:08:41.850

David Tennenhouse: And to build, you know, all the extra data center space and all the extra HP AC and power conditioning and everything else that would have gone into having many, many more servers. So that's been a really big win for VMware and I'd say for sustainability.

43

00:08:42.900 --> 00:08:57.240

David Tennenhouse: Overall, as we extended into other areas like Storage virtualization, you get a similar benefit right you save all those storage devices. Again, their energy from an operating point of view and the embedded carbon so

44

00:08:57.930 --> 00:09:06.090

David Tennenhouse: We it's been sort of really great you know when when both of business for VMware well also contributing to sustainability.

45

00:09:07.710 --> 00:09:21.390

David Tennenhouse: We have, as I said, a sort of even broader not just a commitment to sustainability, but a commitment to being a force for good. So something I do want to emphasize is our interest in seeing the results of this research dedicated to the public.

46

00:09:22.440 --> 00:09:31.860

David Tennenhouse: You know, if there's great research results that VMware can run with and product eyes, that's great if other parts of the industry going product eyes it as well.

47

00:09:32.640 --> 00:09:52.440

David Tennenhouse: That's even better. That's all downstream. We really want to see this fund fundamental research that can basically rises all boats. So we're very excited about this to provide a little more background here, you know, there's a history in the space where

48

00:09:53.580 --> 00:10:03.180

David Tennenhouse: LB and l are kind of the, the folks that produce the regular data center energy reports and starting, you know, in the early 2000s, of kind of when the first least started coming out.

49

00:10:03.840 --> 00:10:14.640

David Tennenhouse: And there was a tremendous amount of banks gnashing of teeth. Because data centers energy consumption was growing at a phenomenal rate. Now we're all kind of actually excited because data centers were growing at a phenomenal rate.

50

00:10:14.880 --> 00:10:29.940

David Tennenhouse: Our field was growing at a phenomenal rate. But meanwhile, our energy consumption, you know, people were drawing these graphs about how soon you know the IT industry would be, you know, consuming as much energy as many countries, etc.

51

00:10:31.170 --> 00:10:38.310

David Tennenhouse: But we gone on to a much better curve, right. So since 2005 the growth rates actually plateaued at a nice rate.

52

00:10:38.730 --> 00:10:45.030

David Tennenhouse: And a lot of the reasons for that worried what I think of as a first generation of innovation. It was actually many innovations.

53

00:10:45.780 --> 00:10:58.650

David Tennenhouse: As I've indicated virtualization was a key contributor to that. And as we virtualize more and more workloads a larger and larger fraction of the workloads in the data center. It sort of helps you know plateau that curve.

54

00:10:59.430 --> 00:11:12.210

David Tennenhouse: Or at least ask them toted somewhat. There is also great work on on redesigning power management at multiple levels. You know, it's sort of in in terms of the servers themselves and being able to

55

00:11:12.990 --> 00:11:22.650

David Tennenhouse: To adjust their power levels and also the power management and power conditioning throughout the data center and then great work on HBC systems, etc.

56

00:11:23.460 --> 00:11:31.920

David Tennenhouse: To improve on the cooling, etc. So we had sort of a, you know, multiple innovations that goddess that photo. That's the great news. It's a wonderful story.

57

00:11:32.250 --> 00:11:47.790

David Tennenhouse: The catchy is that most of that first generation of innovation is itself starting to ask them to there's not that much room to improve the BU ease left and really well run data centers. So we can get more of the data centers to be well run, but we're kind of running out of room there.

58

00:11:48.900 --> 00:11:57.360

David Tennenhouse: We can and are continuing to virtualize more and more of the workloads. We're running out of room there. And meanwhile,

59

00:11:58.110 --> 00:12:04.530

David Tennenhouse: To man for computing is continuing to explode and we have new workloads particularly machine learning.

60

00:12:05.280 --> 00:12:18.930

David Tennenhouse: That you know have a voracious demand right now for computation, which if we don't do something is going to translate into a voracious demand for energy and be a sustainability challenge. So if we could go to the next slide.

61

00:12:21.660 --> 00:12:29.700

David Tennenhouse: I think you know here we really want to see some transformative research that's going to get us, you know, how are we going to keep this

62

00:12:30.390 --> 00:12:43.050

David Tennenhouse: From going from the plateau to another, you know, strong exponential growth. How are we going to ask them to this next round. Right. We don't want to stop people from using more and more computing to do good things. In fact,

63

00:12:44.340 --> 00:12:57.840

David Tennenhouse: A lot of that computation itself goes into improving sustainability of other things. So for example, if people use computation to improve the efficiency of their factory.

64

00:12:58.440 --> 00:13:02.880

David Tennenhouse: Right. They may be dramatically improving the overall sustainability picture.

65

00:13:03.510 --> 00:13:14.730

David Tennenhouse: However, they are increasing the computation demand. And so we want to make sure that can be done in a sustainable way as well. So that's kind of, I think, something we as a field need to bite off and chew on and do

66

00:13:15.480 --> 00:13:22.200

David Tennenhouse: We are hoping that this program will will lead to some really, you know, new breakthroughs and fundamental approaches.

67

00:13:22.830 --> 00:13:29.970

David Tennenhouse: Also looking at the entire system and looking at the total totality of the digital infrastructure.

68

00:13:30.480 --> 00:13:37.950

David Tennenhouse: So we're looking for people to sort of improve the sustainability of what we, you know, the teams referred to her as the digital infrastructure.

69

00:13:38.910 --> 00:13:55.260

David Tennenhouse: And by that they really are, meaning you know kind of all of the stuff that goes into providing our computational networking and storage base. Now having said that what we're not really looking for in the team is going to get into this more detail in a moment.

70

00:13:56.370 --> 00:14:04.110

David Tennenhouse: Is for people to go redesign the mechanical stuff etc and do a whole bunch of Mechanical Engineering Right. The notion iOS, how can we use

71

00:14:04.320 --> 00:14:13.890

David Tennenhouse: You know all computation, the great stuff that the science research community does to improve the sustainability of that overall digital infrastructure.

72

00:14:14.250 --> 00:14:25.620

David Tennenhouse: So there's a difference here between the infrastructure. We want to impact and the tools we use to impact it, which we hope will come from the size community so

73

00:14:26.460 --> 00:14:36.930

David Tennenhouse: You know with that. I'm hoping we will not just be able to sort of keep the curve level, but actually now start knocking the curve down so

74

00:14:37.560 --> 00:14:56.970

David Tennenhouse: Let's set our goals here, not just to offset the new increase demand for computation that we're seeing from machine learning and potentially from other workloads, but to actually knock the curve down so we get on the path towards carbon neutrality and eventually carbon freedom. So

75

00:14:58.290 --> 00:15:02.670

David Tennenhouse: That's sort of the, you know, our aspiration here from VMware

76

00:15:03.780 --> 00:15:17.190

David Tennenhouse: And what I like to do now is turn things over to Curtis, saying, who's the director for the computer and networking systems or I just think of a CNS division in size.

77

00:15:18.990 --> 00:15:35.580

Gurdip Singh: Thank you, David. So my name is good deep saying, I'm the Director for the Division of computer network systems. I joined this March I am on leave from Syracuse University. So, as you mentioned, this is the third

78

00:15:36.930 --> 00:15:44.880

Gurdip Singh: Partnership with VMware, so I was actually program director here at NSF when the first one was light and was very

79

00:15:45.930 --> 00:16:02.580

Gurdip Singh: honored to be part of that in 2016 so so it's great to see this partnership progress and VMware has been a great partner in the giant research that that that has been accomplished. So, as David just mentioned, you know,

80

00:16:04.140 --> 00:16:23.520

Gurdip Singh: The overall goals of this solicitation so want to re emphasize that this solicitation focuses on software layer sustainability of sustainable digital infrastructure. So the proposals with a narrow focus on non software components and disciplines will be considered out of scope.

81

00:16:24.810 --> 00:16:33.690

Gurdip Singh: So in order to set the objectives for the second generation of research in and sustainable digital infrastructure, it's critical that

82

00:16:34.500 --> 00:16:44.070

Gurdip Singh: New matrix and benchmark be defined first. So these metrics and benchmarks should capture the sustainability challenges of the entire computation.

83

00:16:44.730 --> 00:16:53.670

Gurdip Singh: Going well beyond the scope of the first generation matrix where power, you said effectiveness focused on maximizing justice over efficiency.

84

00:16:54.480 --> 00:17:00.060

Gurdip Singh: So in summary, the next generation of innovations for increased sustainability, that we are looking for.

85

00:17:00.450 --> 00:17:07.830

Gurdip Singh: It should include broader considerations and we have provided a list here. So it's matrix and benchmark for systemic

86

00:17:08.370 --> 00:17:17.220

Gurdip Singh: computational efficiency software block bloat and inefficiency distributed resource allocation capacity building and provisioning

87

00:17:17.700 --> 00:17:33.330

Gurdip Singh: compute, storage network trade offs and placement hardware, software co optimization service level agreement frayed space renewable energy source optimization hardware lifetime optimization and renewable energy driven workload shifting

88

00:17:35.220 --> 00:17:37.560

Gurdip Singh: Next, like this. So,

89

00:17:39.030 --> 00:17:46.080

Gurdip Singh: As you know, as it was brought up both by market and by David early at the central question that we want to address it as follows.

90

00:17:46.470 --> 00:17:52.950

Gurdip Singh: What fundamental and systemic approaches systematic approaches in measurement design, development and management.

91

00:17:53.460 --> 00:18:11.610

Gurdip Singh: Of digital infrastructure resources and workload will enable significant progress towards maximizing sustainability of the digital infrastructure with minimal impact on traditional concerns such as programmer productivity and digital infrastructure performance scalability.

92

00:18:13.350 --> 00:18:25.710

Gurdip Singh: So how do we get to this. How do we answer the central question so so for this. The researchers are those, you know, proposing to this program they should consider fundamental nature of the problem.

93

00:18:26.190 --> 00:18:41.580

Gurdip Singh: The multiple aspects of sustainability and the multitude of competing goals of digital infrastructure management. So these competing goals are going to be in terms of the economics performance efficiency and sustainability.

94

00:18:43.020 --> 00:18:55.020

Gurdip Singh: Next slide please. So as we are looking at these research questions. You know what, there are several dimensions or research vectors that we are interested in

95

00:18:55.530 --> 00:19:05.820

Gurdip Singh: And in particular, there are three that to we have focused on here. So one is matrix and benchmark that capture the multiple aspects of sustainability.

96

00:19:06.390 --> 00:19:21.210

Gurdip Singh: Was methods to measure this matrix with for a wide range of applications and infrastructure architectures. Second, our approaches to incorporate sustainability concern across the full development and operations lifecycle.

97

00:19:22.290 --> 00:19:31.470

Gurdip Singh: And third methods to manage aggregate digital infrastructure ecosystem and workloads, with the goal of maximizing the sustainability matrix.

98

00:19:32.310 --> 00:19:41.250

Gurdip Singh: So these research factors that we have noted in the solicitation. They are illustrative of potential key areas of research, but we are not

99

00:19:41.910 --> 00:19:57.570

Gurdip Singh: They are not meant to be perspective or to limit the nature and scope of your proposed work so researchers, they should feel free to propose solutions that address broader digital infrastructure sustainability challenges in novel and innovative ways

100

00:19:59.880 --> 00:20:15.540

Gurdip Singh: So there is a wealth of emerging techniques that could be harnessed to address these challenges that we, the three dimensions that we listed in the previous slide. So examples of such technique include passive and active measurement and systems.

101

00:20:17.610 --> 00:20:23.730

Gurdip Singh: Sorry passive and active measurements of systems at the local at the system, local and wide area levels.

102

00:20:25.260 --> 00:20:33.750

Gurdip Singh: Matrix aggregation services status statistical analysis and inference modeling and learning techniques to assist automated control.

103

00:20:34.380 --> 00:20:54.660

Gurdip Singh: Complex resource management and optimization process and systems isolation to enable control of wide range of parameters and scale of scale at at scale to include sustainability objectives and technologies of agile development in convergence of development and production environments.

104

00:20:55.980 --> 00:21:02.970

Gurdip Singh: So while these above are just examples of emerging techniques that may be relevant to solving sustainability challenges.

105

00:21:03.450 --> 00:21:21.180

Gurdip Singh: Which should not be considered as preferred or the or an exclusive set so proposals are again and courage to consider any

combination of techniques to support approaches that fundamentally increase the understanding of and solutions for digital infrastructure sustainable.

106

00:21:25.110 --> 00:21:33.240

Gurdip Singh: So what our expectations so so so this illustration is looking for proposals for disruptive innovative approaches and ideas.

107

00:21:33.720 --> 00:21:37.350

Gurdip Singh: That have potential to significantly advance the state of the art.

108

00:21:38.010 --> 00:21:52.560

Gurdip Singh: So again, we are specifically excluding research that is primarily in evolutionary improvements over existing state. So we're not interested in incremental, but we are looking at disruptive innovative approaches.

109

00:21:53.490 --> 00:22:04.680

Gurdip Singh: So this solicitation. Also, you know, seeks proposals tak tak Game Changing systems sustainability approaches proposes are encouraged to design a research plan.

110

00:22:05.070 --> 00:22:12.690

Gurdip Singh: With the potential for significant impact on meaningful sustainability matrix. So, this illustration again focuses on

111

00:22:13.140 --> 00:22:23.730

Gurdip Singh: software layer sustainability. Again, we re, re emphasizing this so proposals which have a narrow focus on non software components and disciplines are considered

112

00:22:24.330 --> 00:22:35.940

Gurdip Singh: Out of schools. So, for example, strictly hardware architectures or power, water cooling, you know, these are some things which David also alluded to. During his presentation.

113

00:22:39.420 --> 00:22:49.350

Gurdip Singh: So in this solicitation. Again, we are expecting the proposals to develop prototypes of their proposed approaches to explore implementation aspects of their designs.

114

00:22:50.010 --> 00:22:59.670

Gurdip Singh: And to empirically demonstrate the effectiveness of their solutions. So we want to emphasize that we do not want the community to reinvent the wheel, where possible,

115

00:23:00.690 --> 00:23:07.320

Gurdip Singh: And to avoid unnecessary reinventions proposals should leverage existing software tools and framework where possible.

116

00:23:08.400 --> 00:23:15.990

Gurdip Singh: Same particular researchers, they are encouraged to use existing infrastructures for deploying and testing prototypes of the book.

117

00:23:16.500 --> 00:23:24.480

Gurdip Singh: And for collecting data to demonstrate the effectiveness of their approach so such infrastructure include but they are not limited to

118

00:23:25.080 --> 00:23:41.130

Gurdip Singh: NSF future cloud projects chameleon and cloud lab and also the NSF funded cloud back. So all of these provide NSF funded researchers and educators with access to commercial cloud computing platforms.

119

00:23:42.480 --> 00:23:50.940

Gurdip Singh: However, if the proposal find that the existing frameworks have fundamental shortcomings that cannot be overcome by

120

00:23:51.870 --> 00:24:07.920

Gurdip Singh: That cannot be overcome by a buried from from these existing resources, then the proposals are welcome to develop new frameworks and but they should pay fully explain what the gap is and why do they need to do this entirely new infrastructure.

121

00:24:11.400 --> 00:24:25.260

Gurdip Singh: So as, as you know, you know, we always looking for intellectual merit, as well as broader impacts so proposal should strive to continue to be concrete about the plans to achieve broader impact.

122

00:24:26.370 --> 00:24:39.150

Gurdip Singh: Activities. So these are you know impact of any foundational research of particular interest, our plans for disseminating research results, including open source, open source software.

123

00:24:39.660 --> 00:24:47.220

Gurdip Singh: Production and publication of data sets and activities, leading to real world experimentation measurement and departments.

124

00:24:47.610 --> 00:24:57.660

Gurdip Singh: So wherever appropriate proposals should be clear as to how they plan to navigate God or in broader policy, economic and social considerations.

125

00:24:58.110 --> 00:25:09.120

Gurdip Singh: Which may influence drive or limit the impact of their research proposals many to identify team members with specific expertise to navigate these courses.

126

00:25:10.470 --> 00:25:17.400

Gurdip Singh: So you know so careful attention must replace to, you know, the broader impact of all of the work that that is being

127

00:25:19.470 --> 00:25:37.260

Gurdip Singh: So, so with this. Now I'm next going to turn it over to Matt mode car, who's a program director at CNS, he's going to down for the delve into the solicitation in terms of the requirement and the review process so met mystical

128

00:25:39.570 --> 00:25:54.120

Matt Mutka: And thanks dirty. So I'll just walk you through the proposal submission and review process your work process. And we'll talk about the management of award awarded projects. And at the end, we'll have questions and answers.

129

00:25:57.090 --> 00:26:03.870

Matt Mutka: So, start by noting that the information that you should have immediately. So this program solicitation is NSF

130

00:26:06.030 --> 00:26:14.220

Matt Mutka: Which you can easily find a web search, and it's also posted on the sys website. And as with any to

131

00:26:14.670 --> 00:26:26.880

Matt Mutka: NSF solicitation, you should read it carefully as requirements may be different than from other solicitations and also note that the deadline is 5pm your local time. On November, the fourth 2020

132

00:26:27.780 --> 00:26:34.530

Matt Mutka: submit them to your sponsored research office and plenty of time just to catch the mistakes that are to meet the deadline.

133

00:26:35.490 --> 00:26:48.510

Matt Mutka: An NSF and VMware together. Expect to award approximately two projects and each project may request up to \$3 million over three years and we expect towards to be made in early 2021

134

00:26:50.880 --> 00:26:51.840

Matt Mutka: Proposals may

135

00:26:53.190 --> 00:27:01.140

Matt Mutka: may only be submitted by institutes of higher education and these include universities and two and four year colleges as well as community colleges.

136

00:27:01.620 --> 00:27:17.670

Matt Mutka: Proposals may include an international branch campus of US institution of higher education, including those through several awards and consulting arrangements, but only if the project activities cannot be performed at the US.

137

00:27:22.650 --> 00:27:29.640

Matt Mutka: Individually participate SPI KPI or senior personnel, no more than one proposal submitted in spots to the solicitation.

138

00:27:30.090 --> 00:27:39.120

Matt Mutka: If any individual is found on two proposals, the proposal with the earliest submission date will be accepted and all others will be returned without review.

139

00:27:39.900 --> 00:27:58.170

Matt Mutka: Make sure that individuals on your team have consented to participate on a project and are not listed on another proposal. So we're looking for manageable size teams with individuals with the demonstrate with demonstrated expertise that is needed to successfully conduct the research.

140

00:28:00.270 --> 00:28:11.430

Matt Mutka: So these are in a separate search grants. So we would expect to see graduate students funded as a part of the project. And you may also include software engineers and programmers on solicitation as needed.

141

00:28:12.480 --> 00:28:23.880

Matt Mutka: And note that you have up to 20 pages for the project description. And as always, a check for font spacing margin limitations compliance issues with a proposal.

142

00:28:26.850 --> 00:28:34.200

Matt Mutka: So certain things you should not forget the required postdoc mentoring plan. If you are funding a postdoc

143

00:28:34.770 --> 00:28:42.150

Matt Mutka: And you have up to two pages for the collaboration plan which you must include in proposal if there's one of the one researcher on the project.

144

00:28:42.780 --> 00:28:52.140

Matt Mutka: This collaboration plan is important documented in that notes the appropriateness of each member of the team including the areas of expertise and the role of the research.

145

00:28:52.920 --> 00:29:03.780

Matt Mutka: Document also describes overall management and coordination mechanisms and interdisciplinary dependencies among the tasks which help clarify how the tasks hang together.

146

00:29:04.470 --> 00:29:12.270

Matt Mutka: And finally, there should be a reference to budget items that support collaboration. We expect you to include funds to make the collaboration real

147

00:29:13.830 --> 00:29:20.460

Matt Mutka: And please read the solicitation for more details of what is expected to appear in the document.

148

00:29:21.630 --> 00:29:36.630

Matt Mutka: You may not include letters of general support from for the project or sex or reports expert expected outcomes, but you may include letters of collaboration legend, for example, data or access to resources or facilities.

149

00:29:40.380 --> 00:29:45.180

Matt Mutka: This program will operate under a public dedication intellectual property model.

150

00:29:45.630 --> 00:30:02.400

Matt Mutka: Importantly, the, the solicitation specifies which open source licenses meet the program objectives and solicitation states are requirement that the research results be publicly disseminated through an access compliance repository and in accordance with the NSF Public Access Policy.

151

00:30:05.430 --> 00:30:13.560

Matt Mutka: I like to emphasize that we follow the NSF process of merit review for proposals received to this solicitation.

152

00:30:14.580 --> 00:30:23.130

Matt Mutka: proposals will be reviewed by America review panel and the program directors may request additional ad hoc reviews for a project as appropriate.

153

00:30:24.570 --> 00:30:36.900

Matt Mutka: proposals will be reviewed. According to the intellectual merit and broader impact review criteria and addition reviewers will be asked to considered solicitation specific review criteria found in the solicitation.

154

00:30:39.060 --> 00:30:45.780

Matt Mutka: Ethical Conduct the panel being what Tim key members will attend the panels as observers.

155

00:30:46.470 --> 00:30:57.240

Matt Mutka: And I step in being well where we'll meet to discuss the proposal proposals and decide whether to have a virtual reverse side visit for the top rated projects.

156

00:30:58.230 --> 00:31:09.540

Matt Mutka: And reverse site visit individual teams are called upon to answer questions and concerns regarding raised by the reviews the NSF program directors or the VMware partners.

157

00:31:10.920 --> 00:31:15.600

Matt Mutka: And SF and VMware jointly will make the final funding decisions.

158

00:31:18.960 --> 00:31:26.640

Matt Mutka: The solicitation specific review criteria include the following is going to be the extent to which the project.

159

00:31:27.270 --> 00:31:37.530

Matt Mutka: Has a systems perspective and result in a coral hearing whole that is more than the sum of the parts. We do not want to see individual research topics stapled together.

160

00:31:38.490 --> 00:31:46.110

Matt Mutka: We'd like to see include an evaluation through demonstrations and prototypes of components and eventually as system level.

161

00:31:48.390 --> 00:32:06.750

Matt Mutka: Feature well integrated team of researchers and demonstrate contract concrete plans to impact the broader industry and use existing components and infrastructure where appropriate. Do not reinvent wheel justify your choices, especially when developing new infrastructure.

162

00:32:13.260 --> 00:32:30.510

Matt Mutka: Projects will be funded jointly by NSF and VMware through separate NSF and VMware funding instruments project may receive supplements from either. And so for VMware and NSF expects that the entire fee of all awards will be used to support the proposed work.

163

00:32:33.150 --> 00:32:44.220

Matt Mutka: And SF and VMware will work together to join me manage the NSF VMware partnership projects, while following guidelines and regulations governing each party.

164

00:32:45.360 --> 00:32:49.560

Matt Mutka: In a second VMware will designate a program director for each

165

00:32:50.700 --> 00:32:57.720

Matt Mutka: Ship or the program directors will join or see the execution of the project.

166

00:32:58.350 --> 00:33:06.060

Matt Mutka: VMware will work with NSF to be an active participants, such as advisory researchers on technical issues and industry context.

167

00:33:06.390 --> 00:33:23.490

Matt Mutka: With NSF VMware will conduct annual onsite reviews and may invite attendees to site visits at VMware accepting invitations to attend site visits at at the award. He Ortiz institutions and potentially hosting student interns.

168

00:33:24.570 --> 00:33:35.760

Matt Mutka: VMware may separately fund its own personality to directly support the NSF VMware partnership research either part time or full time with the institution or projects.

169

00:33:40.800 --> 00:33:58.650

Matt Mutka: Note that there are requirements for the project title and I've also listed the expected supplementary documents. Please read the

solicitation carefully and take a look at NSF patchy 2024 overall guidance on writing NSF proposals.

170

00:34:01.590 --> 00:34:06.000

Matt Mutka: Again, the deadline by PM on senators time on November, the fourth

171

00:34:08.760 --> 00:34:16.380

Matt Mutka: And now you can submit questions to the Q AMP a tab and we will

172

00:34:17.670 --> 00:34:19.140

Matt Mutka: Make every effort to answer.

173

00:34:33.960 --> 00:34:34.770

Matt Mutka: Okay, so

174

00:34:36.210 --> 00:34:38.250

Matt Mutka: A big question here.

175

00:34:39.780 --> 00:34:47.730

Matt Mutka: I'm wondering if optical technology such as silicon photonics and optical switching is of interest to this program. Thanks.

176

00:34:49.320 --> 00:34:51.870

Matt Mutka: And Chris, could you answer this

177

00:34:57.720 --> 00:34:58.470

Chris Ramming: I'm sure

178

00:35:01.620 --> 00:35:02.220

Chris Ramming: I'm

179

00:35:04.800 --> 00:35:05.100

Chris Ramming: Pretty

180

00:35:08.400 --> 00:35:09.570

Victor Firoiu: Can take it, Chris.

181

00:35:09.810 --> 00:35:10.620

Chris Ramming: Yeah, go ahead.

182

00:35:11.280 --> 00:35:25.620

Victor Firoiu: Sure. Um, so, as we mentioned in this presentation, the solicitation focus exclusively on hardware specific approaches would be out of scope.

183

00:35:29.910 --> 00:35:35.340

Victor Firoiu: So that includes optical technologies on photonics and optical switching

184

00:35:37.860 --> 00:35:38.580

Matt Mutka: Okay.

185

00:35:41.670 --> 00:35:51.090

Matt Mutka: Given the size of the award is a prop proposal expected to target more than one research vector

186

00:35:52.170 --> 00:36:08.880

Matt Mutka: And with that, I, I would say that we certainly leave it up to the proposers to decide which which or if which one or more than one research factor to do, and given the opportunity you have within the budget.

187

00:36:17.190 --> 00:36:31.530

Matt Mutka: So the next question. Our industry academic partnerships and scope. If so, we assume that confidential information disclose by industry partners will stay confidential.

188

00:36:32.970 --> 00:36:34.140

Matt Mutka: And would

189

00:36:35.520 --> 00:36:38.160

Matt Mutka: Her deep. Why do you want to answer that.

190

00:36:40.260 --> 00:36:42.570

Gurdip Singh: This has his hand up. So if he wants to

191

00:36:42.990 --> 00:37:00.030

Chris Ramming: Oh yeah, I'd love it if you answer it. Um, talk a little bit about the confidentiality of NSF proposal information, but broadly speaking, in order to increase impact of this program certainly industry academic partnerships are welcome.

192

00:37:01.710 --> 00:37:09.540

Gurdip Singh: Right so industry. Yes, as, as Chris said that the industry academic partnerships are welcome and

193

00:37:10.590 --> 00:37:21.660

Gurdip Singh: In terms of confidentiality information so we expect that all of the results that are coming out of this is going to be published and are going to be made open

194

00:37:22.440 --> 00:37:36.990

Gurdip Singh: So, so that's if you read through the solicitation, as well as the so that's what is expected out there on the in terms of the agreement which the part which the

195

00:37:39.930 --> 00:37:50.760

Gurdip Singh: University may have with an industry or partner if they have one on one agreement in terms of how they are maintaining the confidentiality yet.

196

00:37:51.240 --> 00:38:08.490

Gurdip Singh: Trying to also make sure that all of the research that's done under this is is made public that you have to sort of navigate on there, but from the NSF side, we are expecting that all of the results that come out of this funded work are going to be made public.

197

00:38:11.550 --> 00:38:34.620

Matt Mutka: And next question our proposals expected to include or address building a user community for production infrastructure, based on the software layer components. So how are the infrastructure CPU, GPU storage server supports server components needed to do this expected to be provided.

198

00:38:37.800 --> 00:38:42.720

Matt Mutka: I guess I would, I will answer and let someone else all then add add to it.

199

00:38:43.860 --> 00:38:55.260

Matt Mutka: So I mean each proposer is to decide which part that they want to build and if the user community is one aspect of what you would like to have a broader impacts.

200

00:38:55.560 --> 00:39:10.860

Matt Mutka: That certainly is something which you can do who we are just not laying out in the expectations of what we expect to see on that other than to meet the research requirement of the post proposal solicitation.

201

00:39:16.740 --> 00:39:21.180

Gurdip Singh: Yeah, then many of these, I just want to add to these, you know, you sort of need to see

202

00:39:23.130 --> 00:39:32.880

Gurdip Singh: The looking at both intellectual merit and broader impacts, you know, and this clearly some of this comes under the broader impacts categories building the community and you have to

203

00:39:33.450 --> 00:39:42.480

Gurdip Singh: See within the scope of the project and the budget, what are the set of activities that you can meaningfully meaningfully engage and based upon that.

204

00:39:43.650 --> 00:39:46.680

Gurdip Singh: Propose propose your activities.

205

00:39:48.990 --> 00:39:58.590

Matt Mutka: Next question is, is use a VMware software components in any way required to respond to this solicitation and I'll just say the answer to that is no.

206

00:39:59.670 --> 00:40:05.220

Matt Mutka: Not required the VMware software components are not required for this solicitation.

207

00:40:10.890 --> 00:40:18.180

Matt Mutka: So here's another question regarding intellectual property. So all intellectual property must be made public. Is this correct to the

208

00:40:19.350 --> 00:40:22.650

Matt Mutka: Mantra UNICEF normally allows grantees to retain

209

00:40:23.490 --> 00:40:37.770

Matt Mutka: Principal legal rights intellectual property developed under NSF grants to provide incentives for development as emanation of invention software and publications that can enhance their usefulness accessibility and upkeep, what is the motivation for this change.

210

00:40:39.450 --> 00:40:43.260

Matt Mutka: I mean, I think you mentioned the last one, would you want to respond to this.

211

00:40:44.430 --> 00:40:54.600

Gurdip Singh: Yes, so. So again, we are not changing the NSF policy we are trying to say. Only thing I was trying to mention in that. Yeah, the IP is is whatever the normal

212

00:40:55.470 --> 00:41:11.400

Gurdip Singh: Policy that we have where where the universities do retain the right to any kind of inventions all trying to say is that the results on the publications that are from the research results that are expected. They are, they should be in public domain.

213

00:41:11.700 --> 00:41:12.300

Chris Ramming: As well.

214

00:41:12.450 --> 00:41:15.540

Gurdip Singh: So it is no different from any of the other standard kinds

215

00:41:16.260 --> 00:41:25.620

Chris Ramming: I'm actually pretty, pretty sure about that. We had actually in order to achieve broader dissemination of the results we did intend

216

00:41:26.910 --> 00:41:29.790

Chris Ramming: That the research results be dedicated to the public.

217

00:41:32.670 --> 00:41:35.640

Chris Ramming: So there's a special section in the solicitation. That's detail.

218

00:41:39.270 --> 00:41:39.690

Matt Mutka: And

219

00:41:39.900 --> 00:41:43.920

Matt Mutka: We refer to the solicitation and that's what we will move to

220

00:41:49.170 --> 00:42:02.370

Matt Mutka: The next question for climate or proposal are expecting each institution submit separately or should this be structured as several more, you can have a submit separately as collaborative proposals, they do not need to be structured as sub awards.

221

00:42:10.500 --> 00:42:18.870

Matt Mutka: Question. What's with the size. The award is it mainly targeted at large project teams across multiple disciplines and institutions.

222

00:42:21.900 --> 00:42:24.270

Matt Mutka: Any director, would you like to respond to that.

223

00:42:25.380 --> 00:42:29.010

Victor Firoiu: This is not expressly inquired would be

224

00:42:30.150 --> 00:42:33.630

Victor Firoiu: Up to the proposers obviously the

225

00:42:35.490 --> 00:42:36.750

Victor Firoiu: Work proposed.

226

00:42:38.670 --> 00:42:40.410

Victor Firoiu: Would be commensurate to the

227

00:42:41.790 --> 00:42:42.390

Victor Firoiu: Value.

228

00:42:43.440 --> 00:42:51.060

Victor Firoiu: So no, it is not required for report them. It is up. The key is

229

00:42:53.580 --> 00:42:55.830

Matt Mutka: OK, question about

230

00:42:57.660 --> 00:43:02.250

Matt Mutka: How important is to have VMware collaborator VMware collaborators on the

231

00:43:02.250 --> 00:43:08.850

Matt Mutka: Project they be a conflict of interest in the review process. So you wouldn't

232

00:43:09.870 --> 00:43:21.060

Matt Mutka: I will say that any collaborate. You have would be a conflict of interest in review process and and VMware will not themselves personnel will not be collaborators.

233

00:43:25.740 --> 00:43:35.070

Chris Ramming: You know, just come any there are two questions in here that relates to fairness during the proposal process. And I think it's important. You know, I think we intend to release information.

234

00:43:35.670 --> 00:43:45.090

Chris Ramming: Equally and and would not provide any one team with preferential collaboration. So just to be clear on that.

235

00:43:48.780 --> 00:44:02.610

Matt Mutka: Question environmental sustainability for digital infrastructure includes not only research on computing issues, but can also include other areas, such as power read lifecycle and policy or their expectation limits on the balance of the proposed project.

236

00:44:03.720 --> 00:44:16.020

Matt Mutka: I will start by just saying that the expectation is that it is centric to software or for the research problems if the if the

237

00:44:16.770 --> 00:44:33.060

Matt Mutka: Project mostly covers such things as power grid other areas, then it's considered out of scope these other areas can be part of the problem, or what you're considering, but the emphasis and the primary objective on are looking at

238

00:44:33.720 --> 00:44:40.230

Matt Mutka: Innovations for in software for dealing with environmental sustainability.

239

00:44:55.170 --> 00:44:59.970

Matt Mutka: Was a question is the focus mainly on traditional data centers or

240

00:45:00.990 --> 00:45:06.150

Matt Mutka: Wounds also sustainability of cloud edge infrastructure.

241

00:45:07.650 --> 00:45:08.550

Matt Mutka: Certainly.

242

00:45:11.370 --> 00:45:21.870

Matt Mutka: Can you look at what ours, you know, through the software stack what our innovations, you can have to move sent standard building so

243

00:45:22.230 --> 00:45:32.310

Matt Mutka: Cloud edge infrastructures like League would be a place that were some things can do with your choice to decide what what items that you would like to emphasize

244

00:45:42.120 --> 00:45:47.640

Matt Mutka: In the question given sustainability competing is the focus of CNS or

245

00:45:49.110 --> 00:45:55.140

Matt Mutka: Programs. Can you describe a project in scope or this call, but not in scope for

246

00:45:56.550 --> 00:45:58.470

Matt Mutka: Core programs.

247

00:46:01.140 --> 00:46:02.700

Matt Mutka: I think certainly

248

00:46:04.230 --> 00:46:10.770

Matt Mutka: Some of the problems of which you may want to submit or answering this call could also be considered.

249

00:46:11.880 --> 00:46:12.420

Matt Mutka: Core

250

00:46:13.500 --> 00:46:32.040

Matt Mutka: Projects. On the other hand, this is the priests projects can go up to \$3 million over three years, which is certainly not something that which the core support through its medium or large programs. So that would be something which should be something

251

00:46:33.930 --> 00:46:39.000

Matt Mutka: That would be in the scope of such a program sizes that

252

00:46:45.180 --> 00:46:45.540

Matt Mutka: See

253

00:46:55.440 --> 00:47:03.270

Matt Mutka: One asked, could we please get send us a copy of the attendee list so we can have a potential collaboration opportunities things

254

00:47:03.810 --> 00:47:13.350

Matt Mutka: I don't think that's something that we're could could provide out and and and some of the attendees may come and as anonymous. So I don't think we can do that.

255

00:47:17.730 --> 00:47:18.840

Matt Mutka: Seems so

256

00:47:20.100 --> 00:47:21.570

Matt Mutka: One other one

257

00:47:22.740 --> 00:47:41.670

Matt Mutka: Follow up on the optics question, can we explore new optical hardware based design components and their integration with software to uptime optimized sustainability. We can we explore how existing optics based VC cluster designs can leverage software help to improve sustainability.

258

00:47:43.080 --> 00:47:46.800

Matt Mutka: Go ask again, perhaps, Victor, would you like to respond.

259

00:47:47.280 --> 00:47:49.800

Victor Firoiu: Yes. And thank you for that question. This is

260

00:47:51.840 --> 00:47:53.400

Victor Firoiu: Good clarification to make

261

00:47:54.450 --> 00:47:55.200

Victor Firoiu: While

262

00:47:56.220 --> 00:48:09.900

Victor Firoiu: Research and advances in hardware, including optics is not in scope exploiting advances by heart by software is in scope. So

263

00:48:16.110 --> 00:48:21.510

Victor Firoiu: If by exploring new design means

264

00:48:24.660 --> 00:48:37.590

Victor Firoiu: exploring ways software based ways to exploit advances in hardware, then yes, that is in scope, exploring the design of hardware is nothing school

265

00:48:40.680 --> 00:48:41.040

Matt Mutka: OK.

266

00:48:46.980 --> 00:48:56.430

Matt Mutka: OK. So I have a question here regarding says so section 730 of NSF to

267

00:48:57.810 --> 00:49:08.700

Matt Mutka: Live with thousand to Chapter seven has contradicts everything is public aspect of solicitation si, si chapter 18 title varies five of the

268

00:49:10.020 --> 00:49:25.710

Matt Mutka: Common call the bagel act, I will ask us to get the further clarification on anything and make this answer available should be unavailable when we put up all of the information regarding set this

269

00:49:26.790 --> 00:49:32.880

Matt Mutka: Webinar. So that'll be something which I think will have to come back to later unless you're Jeepers

270

00:49:34.260 --> 00:49:48.720

Gurdip Singh: Yeah, I think, yeah, we as, as Matt said we'll come back to you on this one but but as far as this particular solicitation is concerned dia it's specifically indicates that everything

271

00:49:49.590 --> 00:49:55.260

Gurdip Singh: That all of the awardees will agree to dedicate to the public. All intellectual property.

272

00:49:55.890 --> 00:50:09.780

Gurdip Singh: Resulting from the research funded under these awards. So that's what the award conditions are as far as this particular solicitation is concerned, and we will have further clarification on the question that was supposed to

273

00:50:12.540 --> 00:50:23.550

Matt Mutka: Hey, here's a other question say it's difficult to see environmental sustainability and carbon reduction meaningfully address without significant interaction with heart where infrastructure.

274

00:50:24.330 --> 00:50:38.640

Matt Mutka: On solicitation is focused on software components. How much coupling Parker and Cole and optimization can be including without crossing over into our where's the sake airport

275

00:50:41.610 --> 00:50:46.140

Matt Mutka: Perhaps I don't know Chris with, would you like to respond to that or

276

00:50:50.670 --> 00:50:52.230

Chris Ramming: If it's a tricky question.

277

00:50:53.550 --> 00:50:56.490

Chris Ramming: You know, leveraging advances in hardware.

278

00:50:57.720 --> 00:51:01.380

Chris Ramming: In order to have better software control is a great opportunity.

279

00:51:03.840 --> 00:51:11.910

Chris Ramming: I think, I think the intent here was not to have proposals that are focused on hardware related research.

280

00:51:12.570 --> 00:51:27.660

Chris Ramming: And hardware specific advances. So we believe there are a lot of opportunities at the software layer and really our goal is to focus on software layer control as opposed to new hardware developments, so there can be some coupling, how much

281

00:51:29.340 --> 00:51:35.040

Chris Ramming: Is is a is a judgment call. But the intent of the solicitation is to focus on the software layer.

282

00:51:37.470 --> 00:51:38.490

Matt Mutka: A and

283

00:51:39.690 --> 00:51:43.440

Matt Mutka: Another question is the number of awards strictly limited to two or

284

00:51:44.970 --> 00:51:56.280

Matt Mutka: Is it more than two depending on requested budget from attacks. Um, I guess I would answer that the expectation, it'll be too. It's the expectation is what we think, but

285

00:51:58.080 --> 00:52:04.560

Matt Mutka: Other than that, for deep, would you say anything other than that expectation is to and that's what we would

286

00:52:05.130 --> 00:52:12.990

Gurdip Singh: Think yes that's what the expectation is, and clearly if it was the case that there were three that were

287

00:52:14.880 --> 00:52:23.790

Gurdip Singh: Often pressed or maybe there is only one that is of interest, so things can change based upon availability of funds but that's what the expectation is right.

288

00:52:25.770 --> 00:52:26.310

Matt Mutka: And

289

00:52:28.200 --> 00:52:39.270

Matt Mutka: No says will you post the webinars live and my expectation is that something that we normally do. So I, I believe they will appear at some point coming up.

290

00:52:46.320 --> 00:52:48.270

Matt Mutka: Another question coming in.

291

00:52:49.410 --> 00:52:54.900

Matt Mutka: Is, are there any guidelines on experimental platform to demonstrate the value propose.

292

00:52:55.620 --> 00:53:06.060

Matt Mutka: Ideas, can they be specific to propose project can use resource P is have access to her expectations of the part of the proposed ideas that should be demonstrated on a common experiment specified in the cloud.

293

00:53:07.980 --> 00:53:13.980

Matt Mutka: You guys provide letters from other NSF projects is required for is for resource access

294

00:53:16.500 --> 00:53:17.010

Matt Mutka: So,

295

00:53:22.470 --> 00:53:37.830

Chris Ramming: I mentioned that the solicitation, you know, encourages people to use NSF sponsored cloud test beds and there are several that are mentioned in the solicitation. You know, we at VMware hope to provide a limited amount of internal experimental

296

00:53:39.390 --> 00:54:01.170

Chris Ramming: cloud resources and will determine what we can do at a later time. We also may be able to provide some measurements data can be useful. So meanwhile, we're hoping that the teams have some empirical data of their own and and if they are working with other companies or partners.

297

00:54:03.090 --> 00:54:15.270

Chris Ramming: You know, certainly those letters are are welcome. And, you know, we, I think that's all very much in scope for the intent of the solicitation to have broader impact.

298

00:54:17.640 --> 00:54:27.390

Erik Brunvand: suspect there's a tiny bit of confusion on the slide that said zero letters of general support. So I think the subtlety here is that we

299

00:54:27.780 --> 00:54:34.860

Erik Brunvand: Will we do not want letters that simply say this is a great research topic we support them looking at this

300

00:54:35.610 --> 00:54:52.710

Erik Brunvand: The letters of collaborative support will typically provide some actual tangible support, whether it's person time or resource access or something like that. And those are great, but the letters. We don't want other ones from people who say, this is a great idea. I hope they're successful

301

00:55:07.050 --> 00:55:07.500

Matt Mutka: A

302

00:55:09.300 --> 00:55:11.730

Matt Mutka: Couple more questions coming in.

303

00:55:13.260 --> 00:55:19.680

Matt Mutka: already existing data sets for VMware that P eyes may use to motivate the project or proposed solutions.

304

00:55:20.370 --> 00:55:35.280

Matt Mutka: I think, I think that there's nothing that's to be made, other than what's already public and so I would suspect there's nothing that's out there right now in advance of your solicitation advancing your proposal submission

305

00:55:37.380 --> 00:55:40.080

Chris Ramming: That internally. We've identified a number of opportunities.

306

00:55:42.120 --> 00:55:48.600

Chris Ramming: You know, but we do hope that proposers will will bring forward their own data in there and identify their own opportunities.

307

00:55:49.830 --> 00:55:56.130

Chris Ramming: During the proposal period here and there was a related question about, you know, whether or not

308

00:55:56.640 --> 00:56:09.390

Chris Ramming: Any such data would be released equally to all proposal's if it, if, if, if we were to do so, certainly it would be made available to all of these proposals proposals, but there is not an intention to do so at this time.

309

00:56:16.140 --> 00:56:20.250

Matt Mutka: Here's a question. Learn support from NSF proposals are restricted in the current guidelines.

310

00:56:20.700 --> 00:56:39.000

Matt Mutka: Generic phrasing. Will this proposals allow more detail or support not only mentioned know you're you're only allowed what the patchy require has and basically you can have no letters of support at all, but you can have letters of collaboration as specified in SF Pepsi.

311

00:56:42.360 --> 00:56:42.930

Matt Mutka: And

312

00:56:44.040 --> 00:56:49.680

Matt Mutka: This is a lot of existing solutions have individually focused on different aspects of the call.

313

00:56:50.880 --> 00:57:05.100

Matt Mutka: One thing that's missing is a combined solution across the full stack is that specifically what the call is looking for or example several renewable renewable power solutions have already been proposed and research and even implemented to some extent.

314

00:57:06.660 --> 00:57:10.080

Matt Mutka: I saw micro data center at Rutgers, for example.

315

00:57:11.670 --> 00:57:19.410

Matt Mutka: My I would just say that we're not specifying what your innovation is can

316

00:57:20.670 --> 00:57:27.450

Matt Mutka: Look at s different aspects or across the full stack and be your choice. So what is your

317

00:57:30.900 --> 00:57:32.730

Matt Mutka: Victor, would you like to also answer.

318

00:57:33.630 --> 00:57:48.240

Victor Firoiu: Sure, I just want to draw the attention to one paragraph in the solicitation, where it says the solicitation six proposals that deck deco Game Changing sustainability opportunities. So it's really about

319

00:57:49.860 --> 00:57:52.200

Victor Firoiu: The big impact significant impact.

320

00:57:52.860 --> 00:57:54.090

Victor Firoiu: Something that is

321

00:57:54.720 --> 00:57:58.860

Victor Firoiu: Clearly has not been developed yet.

322

00:58:00.090 --> 00:58:06.600

Victor Firoiu: And whether it is a full stack or kind of two thirds of the stack, it's up to the proposals.

323

00:58:12.480 --> 00:58:13.110

Matt Mutka: Okay, I

324

00:58:13.140 --> 00:58:13.530

Think

325

00:58:15.240 --> 00:58:15.870

Matt Mutka: We have

326

00:58:17.490 --> 00:58:21.300

Matt Mutka: Tried to address all the questions that have come up and

327

00:58:23.280 --> 00:58:28.920

Matt Mutka: With that, I think that we're ready to close the webinar.
Any, any final word.

328

00:58:30.660 --> 00:58:31.650

Matt Mutka: Or deed or

329

00:58:36.750 --> 00:58:48.930

Gurdip Singh: No. So, so, yeah. Just want to thank all the participants
for tuning in and listening to the webinar sly webinar and want to also
thank all of the

330

00:58:49.530 --> 00:59:10.620

Gurdip Singh: Team members from VMware, as well as NSF who made the
presentation as Matt said that we will try to assemble all of the
responses and and come back with any clarification that we have that we
posted on the on the program that page. So thank you once again.

331

00:59:11.310 --> 00:59:12.390

Matt Mutka: Yes, thank you all.

332

00:59:12.870 --> 00:59:13.230

Thank you.

333

00:59:15.810 --> 00:59:16.260

Victor Firoiu: Thank you.