WEBVTT

1
00:00:00.539 --> 00:00:04.140
Manish Parashar: Everybody or Good morning, depending on where you are.

2
00:00:05.850 --> 00:00:13.860
Manish Parashar: My name is manish parish or on the office director for the office with Dan cyberinfrastructure and it's my absolute pleasure.

3
00:00:14.370 --> 00:00:24.840
Manish Parashar: To introduce our distinguished speaker today pete backman Peters co director of the North Western university and argonne national laboratory Institute.

4
00:00:25.410 --> 00:00:36.450
Manish Parashar: for science and engineering, he is a leader in the field and has over his career of 30 years has led different aspects of scientific computing systems.

5
00:00:37.020 --> 00:00:47.130
Manish Parashar: His research is focused on software and architecture for large scale and distributed computing operating systems, as well as networks linking.

6
00:00:48.090 --> 00:00:57.810
Manish Parashar: nsf and other supercomputing centers more recently pete and his team have explored cyber infrastructure for the digital continuum linking.

7
00:00:58.320 --> 00:01:08.460
Manish Parashar: Ai and edge computing for scientific instruments and high performance bidding and cloud computing but he was part of the pioneering array of things project.

8
00:01:09.090 --> 00:01:20.100
Manish Parashar: That launched chicago's edge computing and distributed sensing efforts and, more recently, he leads the nsf mid-scale research infrastructure one stage project.

9
00:01:20.790 --> 00:01:37.860

Manish Parashar: Which is building a nationwide cyber infrastructure for Ai at the edge to support a wide range of applications, including ecological research using nsf neon national ecological Observation Network.

10
00:01:39.570 --> 00:01:42.960
Manish Parashar: Major facility so without delaying any further.

11
00:01:44.400 --> 00:01:47.970
Manish Parashar: Let me turn it over to pete over to you beat.

12
00:01:48.900 --> 00:01:54.720
Peter Beckman: All right, thank you and just double checking you can hear me well and we'll get started.

13
00:01:55.770 --> 00:02:02.340
Peter Beckman: So thank you money shot was very kind introduction and I want to thank also you know amy freelander and.

14
00:02:04.170 --> 00:02:16.170
Peter Beckman: Kevin Thompson who have been encouraging in this project, I will be talking about Ai and what's happening in instrument to computation.

15
00:02:16.740 --> 00:02:25.710
Peter Beckman: And it's really the work of many people, some of whom we will meet today as we look through the slides now I was asked at the very beginning to.

16
00:02:26.400 --> 00:02:39.030
Peter Beckman: To provide a little self introduction, maybe, something that how I got into science or the what was a turning point in my life, and you know, speaking of continuum here from curious kid to curious scientists.

17
00:02:40.200 --> 00:02:52.020
Peter Beckman: I it really wasn't one event but I had fantastic teachers, you know, starting with my parents were both immigrants and teaching me both instruction and technology and other things.

18
00:02:52.590 --> 00:03:03.240
Peter Beckman: To my middle school math teacher and my physics teacher and electronics and then I went to college and you know, Professor Stevens and shaffer and cruickshank and.

19
00:03:04.080 --> 00:03:12.450
Peter Beckman: harbor and all you know contributed so it really was the fact that I had great educators who just kept encouraging me.

20
00:03:13.140 --> 00:03:22.620
Peter Beckman: And that really made it possible, so that I could then go out and do other things and be risky in computer science and I also now with my wife.

21
00:03:23.160 --> 00:03:32.730
Peter Beckman: get to go once a year, and work on technology and other projects, clean water, with a local pastor in Africa, so we get to keep pushing that that forward.

22
00:03:33.480 --> 00:03:46.470
Peter Beckman: So we're talking today about how instruments and computing link and nsf has a amazing set of investments in this area.

23
00:03:47.100 --> 00:03:56.310
Peter Beckman: The ice cube neutrino detector in Antarctica is a perfect example of instrument that then needs computing both with the instrument and back.

24
00:03:56.910 --> 00:04:06.210
Peter Beckman: In the other part of the world and we're connected part of the world, of course, neon that money's mentioned is a enormous instrument.

25
00:04:06.630 --> 00:04:15.360
Peter Beckman: distributed, not just one place but distributed across the United States and it's measuring ecological changes climate change.

26
00:04:15.720 --> 00:04:28.560
Peter Beckman: And it also has this data and processing pattern and, of course, if we look into the heavens even further, we can use the observatory in Chile, as an example where.

27
00:04:29.310 --> 00:04:37.800
Peter Beckman: computation and instrumentation are linked now Those are all examples for the in the in the large I work with a.

28
00:04:38.250 --> 00:04:48.390
Peter Beckman: collaborator northwestern university josiah Hester and even in the small this pattern emerges where lightweight sensing in this case for.

29
00:04:49.230 --> 00:04:59.070
Peter Beckman: covert mask can be powered by movement and breath and and and a small battery and do energy storage and do processing right there.

30
00:04:59.430 --> 00:05:14.400
Peter Beckman: And so, all of these patterns have this common theme where there's some instrumentation there's data that's collected and then we process it whether it's processing by hand like Katherine Johnson.

31
00:05:15.390 --> 00:05:20.580
Peter Beckman: who worked at NASA who process, some of the orbital trajectories.

32
00:05:21.240 --> 00:05:36.840
Peter Beckman: To you know today's modern processing, but this pattern now as we look into the future, where we are now is much more complex the volumes of data are just staggering we know now about big data.

33
00:05:37.770 --> 00:05:51.720
Peter Beckman: And so, in this digital continuum we're no longer waiting and doing the analysis, all the way at the end on a supercomputer, we have to start that analysis right there at the instrument, whether it's an iot like that.

34
00:05:52.920 --> 00:06:03.510
Peter Beckman: Face bit mask of josiah hester's or it's a large telescope or it's an another kind of large instrument and, furthermore, we know from these instruments that.

35

00:06:04.560 --> 00:06:13.230
Peter Beckman: That they produce way more data than we actually send and save and so analyzing that at the full resolution is critical.

36
00:06:14.130 --> 00:06:28.140
Peter Beckman: So that brings us to the State we're in now where we must find a way to bring that kind of advanced processing down to the instrument now in the past for very large instruments.

37
00:06:29.370 --> 00:06:38.610
Peter Beckman: You could design a bespoke configuration that you know fpga a's were used for accelerators and light sources and radio telescopes.

38
00:06:39.030 --> 00:06:48.450
Peter Beckman: To do that processing that initial processing at the edge, but for many instruments we need something that's much more flexible and dynamic.

39
00:06:48.990 --> 00:07:04.440
Peter Beckman: Now you can think of edge computing as kind of you know, almost the opposite of a supercomputer can batch or is as a famous quote, that he said a supercomputer is a device for turning a compute found problem into an io bound problem.

40
00:07:05.520 --> 00:07:20.370
Peter Beckman: And you know for us at the edge it's the opposite and edge computer turns in io bound problem into a compute bound problem, so we start that computation at the edge with lidar hyper spectral images camera.

41
00:07:21.090 --> 00:07:32.970
Peter Beckman: radar instrument towers and then there's also activation and so we're going to inject a new kind of computing infrastructure that allows for activation and.

42
00:07:33.240 --> 00:07:45.390
Peter Beckman: adaptation and then leverage is the tremendous advancements in artificial intelligence, so we can do deep learning inference we could do lightweight training at the edge, we can process and have autonomous action.

43

00:07:45.660 --> 00:07:54.960
Peter Beckman: And on the supercomputer side, we can set up digital twins or do mode ex model driven experiments and process the machine learning.

44
00:07:56.040 --> 00:08:06.060
Peter Beckman: So now I you know I know that it's difficult to move that kind of computation off the edge and there's a part of us that says just by faster networks.

45
00:08:06.900 --> 00:08:16.500
Peter Beckman: And and i've worked on those faster networks before, and so there is a real desire to do that, however, there are really good reasons why we need to move the computation they're.

46
00:08:16.950 --> 00:08:28.890
Peter Beckman: All of these kinds of advanced sensors have more data than bandwidth we just can't move at all so processing the full resolution data is very important there's also times when latency is important, especially for activation.

47
00:08:29.460 --> 00:08:40.680
Peter Beckman: Privacy is another issue, and we, the work we do in the split the city of Chicago means that we don't want to stream video and audio 24 seven from these cameras, we want to do that, processing and then delete it.

48
00:08:41.550 --> 00:08:59.730
Peter Beckman: resilience is another reason, with the data Center there comes a single point of success and being able to distribute processing allows a lot more autonomy and finally there's also a power in battery reasons, where we want quiet energy efficient observation.

49
00:09:00.810 --> 00:09:15.720
Peter Beckman: So let me introduce sage now one of the things that coven taught us is is that people are really the most important part, and so, rather than talking initially about the tech, I thought I just highlight the people.

50
00:09:16.470 --> 00:09:22.710
Peter Beckman: And I know this picture of the team from last summer look strange we don't normally go to work.

51
00:09:23.730 --> 00:09:34.530
Peter Beckman: In that attire although sometimes we do, but last year we were virtual with all of our students, and so we you know snuck out and had a time out.

52
00:09:34.950 --> 00:09:44.850
Peter Beckman: A local river to kayak and get to X outdoors so we could actually talk to students and see students and do some more one on one discussions.

53
00:09:45.660 --> 00:10:03.960
Peter Beckman: But that kind of action is very important than the large distributed project, so the project i'll introduce the team Members very quickly Nicola farriers the deputy director she leads all things Ai for the project Scott Carlos copia works on instrumentation he's a.

54
00:10:05.400 --> 00:10:12.090
Peter Beckman: physicist, but now an atmospheric scientist a Valerie Taylor and Mike pap go work on education.

55
00:10:13.050 --> 00:10:25.380
Peter Beckman: eugene Kelly, formerly worked with neon and now works with me as a collaborator from Colorado state and Raj, who is the chief architect of the node but also co leads the waggle project.

56
00:10:26.340 --> 00:10:36.870
Peter Beckman: lk is a specialist in data but recently she is also become a world renowned leader in fire science and simulation.

57
00:10:37.350 --> 00:10:45.330
Peter Beckman: So we're leveraging that in the project Charlie catlett, who was the original founder of the iot project and manages our urban work.

58
00:10:46.290 --> 00:10:57.720
Peter Beckman: Jim old, who manages risk for the project, but his role is actually much bigger in that he is an expert in biology and you know in ham radio in.

59
00:10:58.260 --> 00:11:19.230

Peter Beckman: computing, and so we leverage that in many places, Dan reed is the chief architect, and of course gives us great direction and how to design a facility like this that can be used Cathy and Helen our project managers and keep us on task, because this is a project.

60
00:11:20.370 --> 00:11:33.240
Peter Beckman: Joe swan tech leads the software team distributed software team and keeps us on an agile framework of rolling out new updates all the time and finally Irene quilters is a.

61
00:11:33.810 --> 00:11:38.730
Peter Beckman: On the Advisory Committee and leads advisory committee, so we go to her and asked her for help.

62
00:11:39.540 --> 00:11:51.090
Peter Beckman: So we'll think Kevin Thompson was a project manager and now dive into the tech part, so there are a lot of hard challenging problems in this space and.

63
00:11:51.570 --> 00:11:54.180
Peter Beckman: I pause here for a minute, this is a project.

64
00:11:54.600 --> 00:12:05.010
Peter Beckman: And projects have milestones and hard problems are things that you accomplish on the way to your milestones so there's a lot of research here that we can't really address, but we.

65
00:12:05.280 --> 00:12:11.730
Peter Beckman: do the best we can, as we deliver on the milestones, but this is a ripe area for computer science for search.

66
00:12:12.300 --> 00:12:21.300
Peter Beckman: This sort of cyber infrastructure for things that the edge is very new territory triggered simulations between edge and cloud that's something that.

67
00:12:21.900 --> 00:12:35.190
Peter Beckman: A postdoc at utah is working on remote instruments have have unique environments and resource management is unique, there really isn't a programming model across the continuum yet and so.

68
00:12:36.210 --> 00:12:44.760
Peter Beckman: it's difficult to to design the pieces of code that are laid out and finally we're in a very unique position these.

69
00:12:46.320 --> 00:12:58.410
Peter Beckman: partners that we're adding a sage to have operational instruments neon, for example, is you know, a 30 year instrument and that data is critical that it's always there and so any.

70
00:12:59.310 --> 00:13:12.240
Peter Beckman: Small additions, we make to add instruments or to do computing has to be done in a do no harm method so that's these are some big challenges and great CS problems that maybe later in the afternoon, we can talk about.

71
00:13:12.960 --> 00:13:22.560
Peter Beckman: So the project builds on the success of what was done with the array of things which really was the first of these kinds of edge computing into an urban.

72
00:13:22.920 --> 00:13:33.570
Peter Beckman: Environment Charlie catlett led that and then we took those concepts and said what would it look like to really make this a facility that could scale up so.

73
00:13:34.380 --> 00:13:42.870
Peter Beckman: We there we go, we started the same project with this concept in mind of a software defined sensor.

74
00:13:43.500 --> 00:13:54.450
Peter Beckman: And I know that's a little strange but i'm going to lay it out here, which is that when you run your software container you know sort of a docker container Ai workflow all the way at the instrument.

75
00:13:54.870 --> 00:14:04.500
Peter Beckman: you're essentially creating a sensor that sensor might determine what a plant species are you seeing or what bird you're listening to or if there's smoke.

76
00:14:04.800 --> 00:14:11.220
Peter Beckman: But the concept here is that, with the right cyber infrastructure, if we can create cyber infrastructure that is.

77
00:14:11.850 --> 00:14:27.360
Peter Beckman: pervasive in instrumentation around the United States then students and postdocs and faculty and scientists could write their own code to live in those instruments and process, they can write software defined sensors.

78
00:14:28.170 --> 00:14:39.240
Peter Beckman: So our goals for sage take that concept of the software defined sensor and say let's let's build a prototype let's build a a reusable infrastructure for that.

79
00:14:39.480 --> 00:14:51.210
Peter Beckman: So, in the case of neon let's put that Ai right in the instrument hut and try to do processing along the path so we're building a new reusable cyber infrastructure.

80
00:14:51.720 --> 00:15:02.820
Peter Beckman: it's got to be very well performing and resilient because people are going to depend on it, it has to use the latest tools docker open CV tensorflow Cooper 90s.

81
00:15:03.150 --> 00:15:09.900
Peter Beckman: And then we have to build a community of people writing code for these components at the edge.

82
00:15:10.620 --> 00:15:20.310
Peter Beckman: And then we of course want to deploy this as an experimental test bed and do some new kinds of computing where we trigger from edge to cloud and cloud to edge.

83
00:15:20.700 --> 00:15:28.500
Peter Beckman: And then train up our students and that gives us the ability to both write the software defined sensors and process data.

84
00:15:28.710 --> 00:15:41.220

Peter Beckman: at a very high rate, so an example here from me on, there are many sensors that sample things like a 40 hertz but yet what is really sent when it's all aggregated is you know, on the Minutes scale.

85
00:15:41.490 --> 00:15:59.520
Peter Beckman: And so, looking for anomalies and looking for interesting data, there is fascinating so we have four key partners, the first three were part of the original proposal, the urban work walk HP ran and we fire in southern California neon and then.

86
00:16:00.630 --> 00:16:11.100
Peter Beckman: We have special thanks to Doug to me and the Oregon hazards lab they have been helping and joining the project, even though they're not officially funded and part of the project.

87
00:16:11.610 --> 00:16:19.920
Peter Beckman: But they find it exciting and we have a student coming from them, the summer and we hope that we can do more work with with them so they're part of the team.

88
00:16:20.460 --> 00:16:32.160
Peter Beckman: So let me describe what is actually being deployed so sage is cyber infrastructure it's designed to run on nodes on things in the.

89
00:16:32.490 --> 00:16:36.720
Peter Beckman: In the environment, now it doesn't have to be ours, we can use a Dell server.

90
00:16:37.080 --> 00:16:47.850
Peter Beckman: So it's cyber infrastructure is software so on the right, you see a server blade that can go in and instrument hot and process that data live and provide those capabilities.

91
00:16:48.300 --> 00:17:01.650
Peter Beckman: On the left is a version that can be outside it doesn't need an instrument hot, it has a cpu gpu and some sensors you know tacked on and it's very expandable you can add any.

92
00:17:02.070 --> 00:17:17.340

Peter Beckman: power over ethernet Center you want just plug it into bottom and and it goes so sage the nodes of sage are one of these styles either it's an x86 server with a gpu or it's an outdoor unit that goes out on a Pole.

93
00:17:17.940 --> 00:17:24.120
Peter Beckman: And the architecture for sage is is easy to describe the implementation is actually.

94
00:17:24.720 --> 00:17:35.040
Peter Beckman: It has taken Joe and the team, a lot of work to do, but the nodes themselves run an embedded Linux that has been strengthen for resilience.

95
00:17:35.550 --> 00:17:42.840
Peter Beckman: It then allows these docker containers and Cooper nettie to run people's code at the edge.

96
00:17:43.140 --> 00:17:51.240
Peter Beckman: And it provides libraries for efficient gpu usage, we have sort of an extreme cybersecurity view the nodes only phone home.

97
00:17:51.510 --> 00:17:58.860
Peter Beckman: there's no way to remotely connect to a node unless it phones home, so there are no open ports is another way to say that.

98
00:17:59.370 --> 00:18:01.470
Peter Beckman: The data is published in the beehive.

99
00:18:01.800 --> 00:18:11.490
Peter Beckman: that's our cloud infrastructure and at the cloud, we have an APP repository so all of those docker containers have to be built, such that they can be.

100
00:18:11.670 --> 00:18:19.380
Peter Beckman: shipped over to the nodes and those containers can then be shared open with all the scientists we're going to get to talk about that.

101

00:18:19.980 --> 00:18:35.580
Peter Beckman: That is a very fun part of the project, so we are up and running, right now, we have nodes spread across the country from the East Texas, all the way up to Alaska to.

102
00:18:36.420 --> 00:18:45.180
Peter Beckman: booty avec Alaska where there's a neon site and there is a sage blade in an instrument hot up there.

103
00:18:45.480 --> 00:18:57.540
Peter Beckman: And so we have a dashboard that lets us look at that, but the most important thing, are the Apps right so in order for this kind of new cyber infrastructure to take hold, we need to.

104
00:18:57.870 --> 00:19:06.960
Peter Beckman: Have the capability to write applications for the edge, that means there's a whole programming environment for what it looks like to write an application that runs at the edge.

105
00:19:07.140 --> 00:19:15.330
Peter Beckman: In the same way that there's a program environment for my phone and I have an APP that can run on my phone my phone is the platform for us.

106
00:19:15.720 --> 00:19:25.830
Peter Beckman: The sage node is the platform and scientists can write code Ai code that runs on the platform and we host those codes in an open.

107
00:19:26.280 --> 00:19:35.820
Peter Beckman: portal and here's the portal, you can go take a look, and it has Apps that we're going to talk about now that run all those Ai components, whether it's.

108
00:19:36.180 --> 00:19:48.090
Peter Beckman: calculating cloud cover or traffic state or analyzing avian bird diversity all of those Apps are things that can be compiled and run and shipped out to the node.

109
00:19:49.080 --> 00:20:05.070

Peter Beckman: Now we've already been deploying in Chicago and so here Sean is standing at the street corner watching one of the first nodes go up the nodes are a little bit heavier than when the array of things project was going, I have one here my.

110
00:20:06.180 --> 00:20:12.930
Peter Beckman: window sill and it's they don't just carry them up, they use they use a pulley and they they tow them up and.

111
00:20:13.440 --> 00:20:30.390
Peter Beckman: We have projects in Chicago like this project on the left is looking at a park where they're modifying how flood how water is absorbed to help reduce flooding, and so the processing that happens at the edge will help us understand what's the environments change to that.

112
00:20:31.710 --> 00:20:43.350
Peter Beckman: We have undergrad researchers who are already using the data from sage to process things and to come up with new algorithms, this is a great bit of work from northern Illinois university.

113
00:20:43.980 --> 00:20:50.370
Peter Beckman: Where a camera on the campus is processing using you know, one of these yellow networks.

114
00:20:51.210 --> 00:20:58.320
Peter Beckman: very straightforward, then they process that they calculate the pedestrians they then plot that and can show.

115
00:20:58.980 --> 00:21:12.120
Peter Beckman: Both of going across the crosswalk to and from, and so they can understand, this is a great student project but also it's very applicable, we can understand how people are using city infrastructure and look at safety issues.

116
00:21:14.010 --> 00:21:20.880
Peter Beckman: One that I mentioned sound event detection, this is another great piece of code and this really demonstrates that.

117

00:21:21.750 --> 00:21:31.230
Peter Beckman: That other people's models like we did with yellow but other people's models can be pulled into the sage framework, so this is a model.

118
00:21:31.620 --> 00:21:40.440
Peter Beckman: yam net, which is produced, I think it's by Google it's available, you can download that model and that deep learning model.

119
00:21:40.980 --> 00:21:53.700
Peter Beckman: Has identification for audio and it buckets things into you know 100 or I guess there's 500 different types of audio that it can identify and process.

120
00:21:54.150 --> 00:22:08.520
Peter Beckman: And so we can take that because we have a programming model for the edge, we can take that rapid in the sage pieces and then put it down on a note and count car Horn honks or sirens or.

121
00:22:09.090 --> 00:22:22.530
Peter Beckman: dog barks or other quality of life, so this is a very important example another one we work with scientists both Chicago and it in Wisconsin on diversity on.

122
00:22:23.070 --> 00:22:27.480
Peter Beckman: species diversity birds are great no pun here canary.

123
00:22:28.020 --> 00:22:37.260
Peter Beckman: About species diversity, because they travel so far, thousands of miles, so when the species shows up and when it leaves is very important to understand the ecosystem.

124
00:22:37.530 --> 00:22:47.370
Peter Beckman: So, again now, this time, using the cornell bird model, we can put that together into the same system and run it here the work of daria of the pen drop.

125
00:22:47.790 --> 00:22:58.500
Peter Beckman: We have a node that is in a prairie in the south, we also did some recording at the morton arboretum and you can see here, looking at the data.

126
00:22:59.130 --> 00:23:08.850
Peter Beckman: You know what kind of birds and what kind of confidence in that bird that the model outputs and so scientists can look through that data and better understand what's happened.

127
00:23:09.960 --> 00:23:13.890
Peter Beckman: The very I mentioned in South Chicago with a paintbrush query.

128
00:23:14.550 --> 00:23:26.220
Peter Beckman: is owned by or it's managed by the nature conservancy and they allowed us to put a sensor there and we are monitoring and understanding bird species again by being able to run code.

129
00:23:26.820 --> 00:23:36.840
Peter Beckman: In the device at the edge a great atmospheric example comes from cloud motion vectors, this is a work of Scott calmness and defender and.

130
00:23:38.010 --> 00:23:45.270
Peter Beckman: Believe it or not, you know, a camera that looks up at the cloud, in this case it's looking down on a reflective Dome that looks up at the clouds.

131
00:23:45.570 --> 00:23:55.200
Peter Beckman: and being able to know in real time where clouds are moving and how and categorize those clouds is an important research topic and again.

132
00:23:55.500 --> 00:24:09.210
Peter Beckman: can't stream all the data process it right at the at the source and here, in this case, there are some hard problems to overcome when ice forms on the Dome that top picture in the middle is ice on the Dome and.

133
00:24:09.480 --> 00:24:20.820
Peter Beckman: Ai has to be able to detect that the bottom picture, of course, when it's raining and it's a foul weather, we also collect training data and fall weather is also can sometimes.

134
00:24:22.230 --> 00:24:39.150

Peter Beckman: lead to foul birds in this case a training data, where a bird landed on the Dome and the Ai has no idea what to do about this, but it's a good example of that anomaly an outlier detection algorithms that that these sort of Ai programs need.

135
00:24:40.500 --> 00:24:53.490
Peter Beckman: of measuring water and snow depth is another key application, there are many, this is a picture from neon there are many sites, though, that put sticks in the ground.

136
00:24:53.970 --> 00:25:02.850
Peter Beckman: To measure snow or water depth and then you know i'll say in its current state of the art is a Grad student or some.

137
00:25:03.510 --> 00:25:14.820
Peter Beckman: Worker would go out and they'd look and they write on the clipboard what that value is, and if you have a camera focus on that you might be able to zoom in enough, and you might be able to look.

138
00:25:15.300 --> 00:25:25.530
Peter Beckman: But really we want to develop Ai algorithms that find that step that find the measurement and use computer vision and here song apart.

139
00:25:26.250 --> 00:25:36.270
Peter Beckman: i've been working on using unity resonate and self supervised learning to identify what the water level is now if we can do that for water and snow.

140
00:25:36.510 --> 00:25:51.420
Peter Beckman: Then all of these cameras that are deployed around with water and stick will let us get live data measurements, where we could put that into hydrological models and find out what's happening is it starting to flood we get really live measurement and.

141
00:25:52.560 --> 00:26:06.060
Peter Beckman: One of our postdocs d'addario is also looking at how to use self supervised segmentation to automatically funding to stick so human doesn't have to go in and and identify the bounding box, and so this is a.

142

00:26:06.780 --> 00:26:21.840
Peter Beckman: merging both what kind of computation we needed the edge with what kind of benefit we're going to get in measuring our environment, whether its diversity in bird species or water depth and snow depth or traffic or pedestrians.

143
00:26:23.190 --> 00:26:36.090
Peter Beckman: Now, one of the projects that Nikola worked on with students is to understand water on this is a park in a economically disadvantaged area of Chicago.

144
00:26:36.990 --> 00:26:44.910
Peter Beckman: That doesn't have a lot of investment in in sewer and water, processing and so there's a lot of places that flood.

145
00:26:45.360 --> 00:26:53.220
Peter Beckman: And so, can we use cameras to design an automatic processing at the edge that would allow us to.

146
00:26:53.580 --> 00:27:03.480
Peter Beckman: make a hybrid logical model that would then know when flooding has happened three hours from now it's likely to flood given all the data that we're collecting in real time.

147
00:27:03.750 --> 00:27:10.110
Peter Beckman: by looking at what's happening on the surface and streets in parks in other places, this is using.

148
00:27:10.650 --> 00:27:22.560
Peter Beckman: Deep lab the two with the resonant one or one backbone and plugging these things together again using the best Ai tools we have with a new kind of environment for the edge.

149
00:27:23.190 --> 00:27:30.660
Peter Beckman: we've also started work with utah and we have a node that's on top of the field museum they're looking down both looking at the mountain.

150
00:27:30.960 --> 00:27:42.960

Peter Beckman: And we have a new kind of air quality sensor that they've attached, but we're also looking with the power project which is looking at new kinds of wireless to connect those.

151
00:27:44.610 --> 00:27:56.430
Peter Beckman: radios with software defined radios to sage node so that you can do processing of what's happening in the spectrum again live in real time, but then open that up to anyone.

152
00:27:56.820 --> 00:28:08.370
Peter Beckman: Who wants to try out their code, with a docker container can ship it over and have it run in the note So these are our examples of things that are happening today with sage.

153
00:28:09.300 --> 00:28:21.720
Peter Beckman: Now we're going to shift into wildfire which has become a very important topic, since the the project started the West has been very dry and.

154
00:28:22.530 --> 00:28:30.660
Peter Beckman: Being able to understand wildfire automatically by using computation running in cameras running.

155
00:28:30.990 --> 00:28:40.710
Peter Beckman: simulation modes on supercomputers and pairing that together across the digital continuum is a great project now Okay, as I mentioned leads that frank Vernon.

156
00:28:41.640 --> 00:28:53.250
Peter Beckman: From ucsd runs HP ran, which is a big networking capability in southern California and Doug to me, has been working with the alert wildfire along with okay.

157
00:28:53.550 --> 00:29:09.600
Peter Beckman: And who is this point hundreds of cameras around the West and can we Ai enable them with edge computing so that we can get better at automatically predicting and understanding smoke so that's the backdrop now let's look at this so.

158
00:29:10.650 --> 00:29:22.740
Peter Beckman: One of the current methods of smoke detection has been just to look at an image and the same San Diego group the California group has a model, it does that it works, it looks but it looks at a single image and.

159
00:29:23.670 --> 00:29:39.630
Peter Beckman: Smoke detection is hard because it's really the motion of the smoke, so there, there are people who are working on lstms and and motion based smoke detection as well in California, in Berkeley and other places.

160
00:29:39.930 --> 00:29:50.010
Peter Beckman: We wanted to also look at could we do self supervised learning, could we use an infrared camera that looks at smoke and understands it's hot and then train.

161
00:29:50.430 --> 00:30:05.400
Peter Beckman: The rgb camera to end the rgb using the rgb camera train a model, so we can look at on the right is the infrared camera and on the left is the real life image and so we started working with.

162
00:30:06.090 --> 00:30:26.250
Peter Beckman: The neon team, and they have a kind of tower called a mvp a mobile deployment platform and eugene Kelly at Colorado state, along with the collaborators romell at neon and others they deployed one of these towers at a.

163
00:30:27.480 --> 00:30:37.740
Peter Beckman: tall grass prairie one of the very few remaining original tall grass perry's in Kansas, and they prepared it and put it there, ahead of a controlled burn.

164
00:30:38.070 --> 00:30:59.280
Peter Beckman: Now, on this Tower are all the neon instruments and then we've augmented those with sage processing nodes we have three sage processing nodes plus one instrument hut, which is a sage blade plus an infrared camera that's been power over ethernet jack N, plus a.

165
00:31:00.750 --> 00:31:08.970
Peter Beckman: A air quality sensor that's been checked in, so now we can look at a full range of information and in a real fire.

166
00:31:09.750 --> 00:31:20.730

Peter Beckman: Understand what's happening, and so this is a video or I just went down with a drone and took a video to minute 45 seconds i'll narrow through this very quickly.

167
00:31:21.450 --> 00:31:31.500
Peter Beckman: So, here are some of the experimental patches have been burned, this is the neon mvp tower, it has an instrument array on the top and as we come down the tower.

168
00:31:32.010 --> 00:31:44.880
Peter Beckman: you'll see, there are three wild sage nodes these have nvidia annex gpus in them and they're processing in real time what it sees the smoke detection algorithms.

169
00:31:45.360 --> 00:32:08.910
Peter Beckman: Also, we have that that sensor so after burning the experimental plots which are the square plots that are in the in the area it's a managed by Kansas state, then they went to start burning the area around the tower and, yes, the the Tower is burned 360 they move the wire so.

170
00:32:09.930 --> 00:32:25.290
Peter Beckman: While there, while they're burning it and they lay patches down now here is some of that video where it shows what's happening on the infrared camera and how we can use that to train what's happening in the visual camera so in a future project we hope.

171
00:32:27.180 --> 00:32:38.190
Peter Beckman: Students and postdocs will take Ai data like that and use it to them train new models that do better at wildfire prediction and wildfire monitoring.

172
00:32:38.670 --> 00:32:46.590
Peter Beckman: And so there's a real human goal to this, we have a technology partner arm has been a big supporter and, of course, all of the.

173
00:32:47.430 --> 00:32:56.280
Peter Beckman: Project folks from both from nsf and from the on and from argon that if they helping drive ways so.

174
00:32:56.970 --> 00:33:17.580

Peter Beckman: That was that tower has now been shut down and the data, though, about 250 gigabytes of data has been harvested and San Diego has already put it up on a site and made a dui so if you're interested in looking at that data and exploring that data from the mtp you can do that.

175
00:33:18.930 --> 00:33:31.830
Peter Beckman: And we also have a collaboration that we're starting now with CSI row and Australia we're sending 10 nodes to Australia, where they will they have a computer science group called.

176
00:33:32.850 --> 00:33:43.320
Peter Beckman: Data 61 and they're looking at new kinds of models again by running code in that infrastructure, they can run models and understand what's happening.

177
00:33:43.590 --> 00:33:51.990
Peter Beckman: At the edge with respect to urban growth water they're also very interesting wildfires, as you know, Australia, has quite an issue with that.

178
00:33:52.500 --> 00:33:59.730
Peter Beckman: Now, at the bottom I showed the user's manual in order to deploy nodes in Australia, without going there, although.

179
00:34:00.390 --> 00:34:12.450
Peter Beckman: it's been made a little bit easier that Kobe has relaxed a bit, we have a very sophisticated installation manual that we can send people who are installing those combined nodes together.

180
00:34:14.040 --> 00:34:32.820
Peter Beckman: We also have a fantastic partnership with the ojibwe nation it's a sovereign tribe in Wisconsin and surrounding states they're very interested in climate change, especially as it relates to my Newman, which is the ojibwe word for wild rice what we call wild rice.

181
00:34:33.900 --> 00:34:44.670
Peter Beckman: And we've been working with Kim from northwestern university and she has been to the tribal Council meeting and prepared and gotten permission from the.

182

00:34:45.240 --> 00:34:53.040
Peter Beckman: From the tribal Council but to allow us to install a node that we can then use to help monitor things like.

183
00:34:54.030 --> 00:35:01.920
Peter Beckman: The river and wild rice and wild rice beds, and so we have the reds over there is where they picked out a place on a Pole it's pretty remote.

184
00:35:02.310 --> 00:35:16.470
Peter Beckman: But we believe we can get enough networking to pull that data in and then work as a partner with the sovereign try to look at the data, so all the data will be hosted and shared with the tribe, and they will have copies of all that data.

185
00:35:17.070 --> 00:35:23.490
Peter Beckman: In a sovereignty agreement that we have with them we're also working with josiah has stirred with a.

186
00:35:24.390 --> 00:35:36.090
Peter Beckman: native hawaiian community and we're hoping to be able to deploy one of the nodes we've just been working on how to get it there and get it deployed, and in this case it's to a traditional fish pond.

187
00:35:37.260 --> 00:35:51.150
Peter Beckman: made by their ancestors, and again they want to study that the impact of climate change and allow students to write code for those those components Now I want to pause here as we're pretty far into the into the discussion.

188
00:35:52.020 --> 00:35:55.980
Peter Beckman: And just point out that we've talked a lot about nodes and so.

189
00:35:56.430 --> 00:36:01.560
Peter Beckman: You might be tempted to start thinking sage is a physical thing, like the know that my desk.

190
00:36:01.830 --> 00:36:11.940

Peter Beckman: But really I want to just remind everyone it's cyber infrastructure it's software that happens to need a physical embodiment in many cases, so we can run.

191
00:36:12.240 --> 00:36:22.290
Peter Beckman: On an x86 blade and use other people's sensors so for the neon example we're pulling data we can pull data from their Fino cams and from.

192
00:36:22.530 --> 00:36:33.120
Peter Beckman: The covariance calculations and pull that right into sage so it's possible, however, when you're deploying outdoors you need to weatherproof that box and that's what's happening.

193
00:36:33.690 --> 00:36:41.760
Peter Beckman: So any kind of sensor can be added, with the exception that we don't do windows, so if you if there's a proprietary windows driver for a sensor.

194
00:36:42.210 --> 00:36:46.500
Peter Beckman: We don't want to work on that, but anything that has open protocol that.

195
00:36:46.980 --> 00:36:58.530
Peter Beckman: Air quality monitor with something that does to me from Oregon recommended we got it and figured out the protocol and set Linux up and it's reporting in, and we can live stream and process, the data in real time.

196
00:36:59.070 --> 00:37:10.650
Peter Beckman: We have a portal that allows us to look at that data in real time so that data is shipped to the beehive and you can zoom in and understand what all the data is.

197
00:37:11.130 --> 00:37:30.150
Peter Beckman: being sent and what's happening to it, it can be connected with a live API So if you have an htc program that you want to pull live data from sage and connect your hbc program to the edge that's possible you can also download you know hard zip file, so, if you like.

198
00:37:31.200 --> 00:37:41.190

Peter Beckman: So we have a lot of online resources from tutorials and the data that was collected the status, you can look at the admin page.

199
00:37:41.670 --> 00:37:48.390
Peter Beckman: we're starting it's not quite released yet, but we're starting to visualize it for the job scheduler so one of the key parts.

200
00:37:49.290 --> 00:37:56.370
Peter Beckman: For this i'll talk a little bit more about is, you have to be able to run shared jobs on a shared infrastructure.

201
00:37:57.030 --> 00:38:06.060
Peter Beckman: kind of like what we do with supercomputers we are sharing the edge and so that's a key feature so we're developing a way to understand visualize that.

202
00:38:06.570 --> 00:38:20.880
Peter Beckman: Now the nodes that are physical our manufacturer, the local manufacturing company and so those are manufactured put on racks and and burned in and.

203
00:38:22.140 --> 00:38:32.400
Peter Beckman: Sit for quite a while as test run and then eventually boxed up and put on pallets and shipped to our collaborators now of course if you're have an instrument hut.

204
00:38:32.790 --> 00:38:40.290
Peter Beckman: Then adele blade or, if you have a big instrument like a telescope then you'd have a rack which you can still run sage with you'd have a rack of servers.

205
00:38:40.620 --> 00:38:54.330
Peter Beckman: And in this case, what you see here is for the manufacturing of those notes the nodes have done been extensively tested so given to a company that freezes them sprays water on them blast them with dust.

206
00:38:54.960 --> 00:39:11.310
Peter Beckman: puts them in a chamber to test rf emissions they, in this case here, they put about five or six inches of ice on it, then let it melt and so we've worked through all of those issues, to make sure that they last and are are in the field, a long time.

207
00:39:12.540 --> 00:39:24.690
Peter Beckman: we've also hosted our first hackathon where we had four people who are connected to the project in some way, whether in Texas and utah in Oregon.

208
00:39:25.470 --> 00:39:37.650
Peter Beckman: To to get on and log into a node and develop an Ai code or try and practice with an Ai code that's running at the edge and reporting back.

209
00:39:37.980 --> 00:39:45.180
Peter Beckman: And so y'all know and Sean love that tutorial effort and we plan on repeating that that's a very important part of our project.

210
00:39:45.750 --> 00:39:58.050
Peter Beckman: Is that that set of applications, we have about 12 applications, now that can be run at the edge and we're talking to partners, all the time about new applications.

211
00:39:58.500 --> 00:40:05.880
Peter Beckman: For example, it's not quite ready that's that's application that measures graduated depth in water or.

212
00:40:06.360 --> 00:40:26.700
Peter Beckman: With snow that's still under development, and we would love for a computer science group at a university to take over a set of applications and develop them and so these hackathons will let us develop the Ai community that can run in that full distributed instrument.

213
00:40:28.050 --> 00:40:33.540
Peter Beckman: Now I I gave a teaser earlier about multi tenancy, this is a hard problem.

214
00:40:34.800 --> 00:40:45.870
Peter Beckman: And it's unlike other problems in that the way we manage most resources sort of supercomputer resources is you submit job to a queue and you wait your turn.

215
00:40:46.800 --> 00:41:05.100

Peter Beckman: And if it happens to me five hours and having these 10 hours you're okay with that you know have other work to do, and you come back and you check your results for the edge, we have a different problem because phenomenon happen in real time it's raining now you know I detect.

216
00:41:06.120 --> 00:41:19.470
Peter Beckman: BATs now i'm interested in this anomaly that's happening in the eddie covariance right now, because the anomalies happening right now and I need that code to be running to process that and.

217
00:41:19.800 --> 00:41:32.610
Peter Beckman: Understand what's happening at the full resolution right now so there's a very difficult almost tense relationship between the urgency of running applications and when they.

218
00:41:33.090 --> 00:41:40.620
Peter Beckman: need to run and how to trigger them so it's a new kind of scheduler and y'all know Kim has been pioneering the work for that.

219
00:41:41.010 --> 00:41:50.310
Peter Beckman: And we've decided that, instead of like time slices like you get a supercomputer, we have to share, we have to multi tenant and, in some cases, you have to share the gpu.

220
00:41:50.700 --> 00:42:00.330
Peter Beckman: And so examples of these kind of science goals might be based on time of day, it might be based on it's hot it's raining it's.

221
00:42:02.040 --> 00:42:12.060
Peter Beckman: The water, the water level is so high, or it might be based on a prediction that's outside of the node, so this is where that model driven.

222
00:42:12.930 --> 00:42:20.670
Peter Beckman: Experiments are so important right the connection across the digital continuum if you have if you know, a hurricane is coming, if you know that.

223
00:42:21.150 --> 00:42:31.650

Peter Beckman: Conditions for flood for tornado or extreme drought are occurring that direct yo that you know event is happening, you should be able to send.

224
00:42:32.310 --> 00:42:42.990
Peter Beckman: code and priorities and set the science goals to the edge, that this is the important thing here's how to set up and run and, of course, that edge component also has a has a.

225
00:42:43.530 --> 00:42:58.410
Peter Beckman: An automation component, so you can see here on the right, we have two different plugins that's what we call the components that run they run and can alternate or they can run simultaneously and, of course, then you know share the resource.

226
00:42:59.520 --> 00:43:16.050
Peter Beckman: So we also have a strong education focus now, this is an area that's actually easy to get students interested in its cyber physical right, you have things that you're looking at their pictures or audio their sensor readings there's temperature there's pressure there's.

227
00:43:17.850 --> 00:43:18.900
Peter Beckman: So many.

228
00:43:19.950 --> 00:43:29.250
Peter Beckman: Things that they're very familiar with that they want to explore and so Mike paprika and Valerie have worked to build a kind, and this was during coven.

229
00:43:30.090 --> 00:43:44.670
Peter Beckman: To build a kind of kit that was cheap 300 bucks and could go to be used by a handful of students includes an nvidia jetson Nano, which is a real you know edge gpu and it allows them to.

230
00:43:45.750 --> 00:43:57.510
Peter Beckman: Do projects and include like, can you tell when your air conditioner is turned on and off by the noise can you run those yellow image detectors what other things, and they all started Jupiter notebooks to do this.

231

00:43:58.020 --> 00:44:07.950
Peter Beckman: So this has been a fantastic project sort of at the at that level but we've also worked all the way down into eighth grade, and we have a.

232
00:44:08.610 --> 00:44:12.180
Peter Beckman: program that has pulled the students in we have helped.

233
00:44:13.140 --> 00:44:25.440
Peter Beckman: With a partnership with and I ustream and with our guns education department to pull in at risk and other students, and so we have a great diverse crowd and we have had several camps that have run.

234
00:44:25.890 --> 00:44:39.960
Peter Beckman: introduction to Ai challenges understanding ethics with Ai and then get to the computation part understanding, even just running the merlin APP to understand what is happening How does that work and what can we write code for to understand.

235
00:44:41.280 --> 00:44:43.500
Peter Beckman: From the edge to the cloud.

236
00:44:44.610 --> 00:44:56.940
Peter Beckman: So what are the next steps for for sage as we look to the future, so we have you know 7080 nodes already in the field, some of them are at the warehouse.

237
00:44:57.300 --> 00:45:08.400
Peter Beckman: During the field burning in the warehouse they're about to be shipped other places nodes are being built we're trying to increase the number of sage blades that are out.

238
00:45:09.390 --> 00:45:18.390
Peter Beckman: with partners, and so our plan over the next six months in a year is to deploy more nodes get more endpoints for edge computing.

239
00:45:19.110 --> 00:45:26.400
Peter Beckman: Also to integrate more advanced sensors now I show a sensor here we've worked with Erin Kennedy before and he's built his own instrument.

240
00:45:26.910 --> 00:45:35.310
Peter Beckman: And that instrument is a snowflake camera there are strobe lights and cameras and as snow falls are many kinds of snow I didn't.

241
00:45:35.730 --> 00:45:50.100
Peter Beckman: Really fully appreciate how many type types of snow and snowflake and you know sleep and other things fall and to do that in real time so that requires real edge processing and visual.

242
00:45:51.060 --> 00:45:59.670
Peter Beckman: Computer vision real time so that eventually you can get real data as to what's falling from the sky in real time.

243
00:46:00.570 --> 00:46:06.000
Peter Beckman: We also want more students more hackathons and and hopefully the summer more kayaking.

244
00:46:06.720 --> 00:46:14.400
Peter Beckman: We will be in person, this summer, and so that is very exciting we've already I saw this morning, one of the students walking down the hall.

245
00:46:14.790 --> 00:46:21.990
Peter Beckman: And so we're just very excited to get them engaged in these Ai Apps we also are looking at more partners.

246
00:46:22.770 --> 00:46:32.880
Peter Beckman: The mount Washington Observatory we've had a wonderful conversation with them and they already have some sensors out there, of course, the problem is that get very, very, very cold.

247
00:46:33.750 --> 00:46:44.610
Peter Beckman: So what we think we'd start with is to put a stage blade and processed see we can process the imagery to give rhyme ice depth.

248

00:46:44.970 --> 00:46:53.220
Peter Beckman: On towers and other things in real time, that would be very exciting that's a future project we're not there yet, but we hope that that's something we could accomplish.

249
00:46:53.700 --> 00:47:07.890
Peter Beckman: there's a great new park in the south side South side Chicago called big marsh it's it's hundreds of acres it's in an old sort of more industrial area but it's a wonderful park where people ride mode ride.

250
00:47:08.700 --> 00:47:15.720
Peter Beckman: bikes and mountain bikes and there's a lot of avian diversity there and they're they're very interested in what.

251
00:47:16.290 --> 00:47:27.180
Peter Beckman: That land use and visitors and what's happening around the city of the nature conservancy we've already deployed with them in the past they're like to expand that the Lincoln park zoo we have a partnership there.

252
00:47:27.600 --> 00:47:40.320
Peter Beckman: And we want to add anomaly detection as well, so what we've been focusing on currently is running code at the edge that is using Ai and deep learning and doing.

253
00:47:41.970 --> 00:47:55.980
Peter Beckman: Essentially, in France, we explored lightweight learning as well, but we noticed that in some of the instrument streams there are anomalies, and this is a real one from the neon mvp.

254
00:47:57.120 --> 00:48:01.860
Peter Beckman: And this is just their their vibration sensor it's a gyroscope it's a you know 3D.

255
00:48:03.120 --> 00:48:11.850
Peter Beckman: You know, vibration right and the neon tower is sampling this data and because they've allowed us to put.

256
00:48:12.270 --> 00:48:23.400

Peter Beckman: The sage infrastructure there to monitor our instruments, but also get live the data stream we look at this data stream and we see that on may 5 there's an event.

257
00:48:24.060 --> 00:48:33.330
Peter Beckman: Now, whether it's any covariance or its vibration these kinds of anomalies are really key to be able to understand okay well what was this event.

258
00:48:33.540 --> 00:48:49.350
Peter Beckman: In this case, we believe it was a workman starting to work on the tower and put new instruments on but that vibration good impact other readings and so all of that kind of anomaly detection across multiple data streams whether it's a telescope or it's.

259
00:48:50.490 --> 00:49:07.680
Peter Beckman: A an instrument that's distributed or for audio processing or hyper spectral are all that kind of anomaly detection in playback so you could pull back the data before and capture the data after an anomaly, is very important for tuning instruments at full resolution.

260
00:49:08.850 --> 00:49:19.080
Peter Beckman: Now, I mentioned that we love students and we have a lot of projects that we would love to see progress on these are things that.

261
00:49:19.590 --> 00:49:33.330
Peter Beckman: While we talk to scientists about the digital continuum they come back to us and say Oh, this is what I would really love, I really love to be able to understand wildlife movement in real time.

262
00:49:34.170 --> 00:49:51.690
Peter Beckman: i'd like to understand bike lane dynamics right How are people using bike lanes or I am I want to understand just using the debris that is floating in a river, I want to understand the speed of the River or looking down through the water it's timidity.

263
00:49:52.740 --> 00:50:07.230
Peter Beckman: or measure calculate biodiversity by listening to me that's again with a project with Susanna in in Wisconsin understanding all whether it's BATs or insects or.

264
00:50:07.830 --> 00:50:21.240

Peter Beckman: mammals or birds measuring biodiversity quickly measuring ice and ice coverage all of these kinds of computations are right for Ai and for students to.

265
00:50:21.630 --> 00:50:31.620
Peter Beckman: bite one of these off to start with one of the initial new networks start with one of the models and then modify it or to just use a new method entirely.

266
00:50:32.520 --> 00:50:40.080
Peter Beckman: And so there is tremendous promise here in motivating and exploring for curious.

267
00:50:40.530 --> 00:50:57.990
Peter Beckman: Scientists and curious students to how they could run code on stage and the great news is they really could run that code on sage they could get an account they could write a piece of code that lives in a docker container and then ship that out and and see the data live.

268
00:50:59.040 --> 00:51:09.600
Peter Beckman: So we're back to this question, which I shared at the beginning, which was, as we look at programming this new digital continual.

269
00:51:11.190 --> 00:51:21.060
Peter Beckman: We have decades of experience on this on the supercomputer side from mpi and math libraries and Patsy and.

270
00:51:21.690 --> 00:51:39.660
Peter Beckman: Job schedulers and the great work of the networking teams that put 100 gigabit links between all of the all of the center's and now places like tanker are injecting Jupiter notebooks and live visualization, and so we have tremendous amount of experience there.

271
00:51:41.040 --> 00:51:44.760
Peter Beckman: What can we do and how open could we make.

272
00:51:46.050 --> 00:51:53.970
Peter Beckman: The edge to link that full digital continuum could we put sage on a radar.

273

00:51:55.050 --> 00:52:04.230
Peter Beckman: truck right these scalable drivable radar trucks, we put them in the planes, can we do live processing on drones.

274
00:52:04.500 --> 00:52:17.280
Peter Beckman: What can we connect telescopes microscopes where will we do, processing and what will that cyber infrastructure look like that's different from our cloud infrastructure that we have with.

275
00:52:17.850 --> 00:52:23.610
Peter Beckman: Amazon and sure and and Google and even our own sort of nsf and.

276
00:52:24.210 --> 00:52:34.830
Peter Beckman: clouds and infrastructure, like chameleon what does that look like, as we pull that down like what sage is doing, can we build a full environment of scientists.

277
00:52:35.490 --> 00:52:44.430
Peter Beckman: That link large scale and SF facilities and just strip it and facilities and allow people to program those facilities.

278
00:52:45.300 --> 00:52:57.390
Peter Beckman: So with that I want to thank there, we have a lot of folks who have been part of the project we've had students since 2013 when we first started exploring.

279
00:52:58.200 --> 00:53:08.010
Peter Beckman: Ai and edge computing at that time that term edge computing wasn't really used so we were struggling with what do we call it, we want to process right there at the at the instrument.

280
00:53:08.610 --> 00:53:17.550
Peter Beckman: And in 2020 we were fully fully virtual but we're very excited about getting back and using those student teams.

281
00:53:17.850 --> 00:53:33.000

Peter Beckman: Throughout the year and working with San Diego and with Colorado Oregon utah Texas, and all of our partners in deploying new kinds of technology to allow new instruments and new science to be done.

282
00:53:34.110 --> 00:53:42.330
Peter Beckman: So with that, I think we have time for questions I last night rajesh found this picture it's from.

283
00:53:43.740 --> 00:53:54.840
Peter Beckman: One of the very, very first installations, of the proto work that would eventually become waggle and sage at the nature conservancy and there.

284
00:53:56.340 --> 00:54:04.410
Peter Beckman: i'm sure they're working safely scrapping this up there and Liliana one of their PhD students was there to snap a picture so.

285
00:54:04.980 --> 00:54:17.190
Peter Beckman: we're we're happy about those from from humble beginnings to now working in a distributed way across the country so with that money shall turn it back to you and I can either stop sharing or.

286
00:54:18.630 --> 00:54:19.680
Peter Beckman: Look at some slides.

287
00:54:20.640 --> 00:54:24.210
Manish Parashar: So thank you pete for an absolutely fascinating talk.

288
00:54:25.710 --> 00:54:28.140
Manish Parashar: A virtual round of applause repeat place.

289
00:54:29.520 --> 00:54:35.400
Manish Parashar: We have a few questions that have come in, so let me read them to you, and maybe we can get you to answer them.

290
00:54:36.870 --> 00:54:47.520

Manish Parashar: So the first one is from Donald walsch What if somebody wants to pay for another stage node to be added at a place of their choosing.

291
00:54:48.840 --> 00:54:54.030
Peter Beckman: that's a that's a very good question and, in fact, our as a program manager.

292
00:54:54.600 --> 00:55:05.790
Peter Beckman: Kevin has been pushing us how that question because that scalability is very important, so we, we can obviously produce these nodes in the electronics company can produce them for us.

293
00:55:06.510 --> 00:55:12.210
Peter Beckman: And the hard part is is integrating them into that full infrastructure and so.

294
00:55:12.810 --> 00:55:26.850
Peter Beckman: We are able to do that and we'd love to talk to you about your science problem, but one of the questions that we've had we've had several nsf projects proposed to deploy this technology as part of their project.

295
00:55:27.570 --> 00:55:40.590
Peter Beckman: So, both for the ojibwe but also in the city of Chicago and also a project and Finland and other places where the next step is to to be able to.

296
00:55:41.640 --> 00:55:52.290
Peter Beckman: make it more available so that people can get the both a sensor either they can buy the the sage blade a Dell server or they can get the.

297
00:55:52.500 --> 00:56:01.980
Peter Beckman: version that goes outdoors and then integrate the software so we're in some sense right now we're think of us as a as a cloud provider, in that case, which is that.

298
00:56:02.460 --> 00:56:20.850
Peter Beckman: We would give an account and we would let people log in to their node and start to develop Ai applications for it so we'd love to talk to you about places where processing data at the edge as possible and we're trying to work out of a more.

299
00:56:22.620 --> 00:56:28.950
Peter Beckman: well structured plan for our external partners who are not sort of nsf partners would be able to do that.

300
00:56:31.380 --> 00:56:43.470
Manish Parashar: And thank you pete so here's another one from Marcus pen while five g and 60 technology improves the bandwidth latency from the edge to edge BC significantly.

301
00:56:44.220 --> 00:56:53.550
Manish Parashar: Will you still recommend to process the same data that you're presented over here what seismic change would you envision for the edge processing.

302
00:56:54.330 --> 00:57:01.530
Peter Beckman: yeah no that's a very good really insightful question I am i'm going to throw up this slide here so.

303
00:57:02.100 --> 00:57:24.600
Peter Beckman: There are some projects that are happening at the D that are using the core waggle platform that's the hardware box and for understanding vehicle understanding nonproliferation in urban and in a credo with excellent, but also with 5g and digital and we work with Randy.

304
00:57:25.890 --> 00:57:45.240
Peter Beckman: Randy Barry at northwestern university and your question is spot on which is one of the key features of five g is that technically you're able to be able to run code in the tower, so what they call the mech right.

305
00:57:46.470 --> 00:57:55.200
Peter Beckman: And the Multi user tower now the way verizon and and other companies are doing this right now is sort of virtual machine.

306
00:57:55.920 --> 00:58:14.160
Peter Beckman: They have a rack of x86 servers in the base of the tower and there are a lot the The hope is that that eventually becomes like a cloud service for 5g providers, and so the idea of course would be if you have a game or APP that needs a low latency.

307
00:58:15.420 --> 00:58:20.610
Peter Beckman: Then you would just connect to the tower and code to run right there to give you, your answer.

308
00:58:21.840 --> 00:58:29.040
Peter Beckman: So that means that this digital continue to answer your question now let me it's a digital continue has multiple sites along the path.

309
00:58:29.400 --> 00:58:38.850
Peter Beckman: So there's the edge edge right all the way at the instrument, then 5g and other new advanced networking and things like fabric.

310
00:58:39.510 --> 00:58:59.400
Peter Beckman: might provide computing one step in into the network and then finally go all the way to the cloud so there's tremendous opportunity five g is going to lower latency is in many cases and give us improve bandwidth terrible bandwidth and in our project we're looking at at.

311
00:59:00.660 --> 00:59:16.410
Peter Beckman: Choosing the spectrum to use kind of like the advanced work powder project is doing, but then there are more locales more places to run code and that will make this programming model that I said, is currently missing, you know even more important for.

312
00:59:16.410 --> 00:59:17.040
us to solve.

313
00:59:18.780 --> 00:59:24.870
Manish Parashar: Thank you so The next question is asking, and this one to Deborah.

314
00:59:28.110 --> 00:59:37.830
Manish Parashar: pinch off it says, thank you pete This is great, is there a mailing list or the communication channel to get announcement for future hackathons.

315
00:59:38.520 --> 00:59:48.300

Peter Beckman: event, yes, that would be great so on the stage website there's a there's a safe continuum.org I believe there's a link to the.

316
00:59:50.580 --> 01:00:01.740
Peter Beckman: The help line I think it's support at sage continuum.org and if you just email that will get you connected up so support at sage continuum.org and that'll get to them.

317
01:00:03.180 --> 01:00:03.660
Manish Parashar: Thank you.

318
01:00:04.710 --> 01:00:21.240
Manish Parashar: there's a question here from Julie bonsai it says, if we have edge computing on multiple facilities, can you elaborate on the challenges of integrating the data to get a holistic perspective on the couple climate system thanks for a great dog.

319
01:00:22.380 --> 01:00:35.280
Peter Beckman: Well yeah that's a that is a hard problem and it's a very good problem is that one of the things that we've we noticed is that most of the sensors.

320
01:00:36.690 --> 01:00:43.380
Peter Beckman: that the different groups put out different sensors so there's the America flux set of sensors there are neon sensors.

321
01:00:44.100 --> 01:00:55.380
Peter Beckman: There are the D has a set of sensors called arm, which are in the southern great plains, we have a in Oklahoma we have a note there, and these things are generally not.

322
01:00:56.130 --> 01:01:15.150
Peter Beckman: Well, connected they each go to their own data bucket in the cloud, and so I think there's tremendous opportunity in two areas, one is how do we fuse that data well for architectures and scientists who want to study climate or.

323
01:01:16.170 --> 01:01:22.410
Peter Beckman: movement or human activity or urbanization, how do we put those two data sets together.

324
01:01:22.920 --> 01:01:34.920
Peter Beckman: But second at the edge, can we be more flexible in gathering data, such that we can share it across multiple instruments across neon across America flux across arm across.

325
01:01:35.790 --> 01:01:52.320
Peter Beckman: The next gen drones that will fly and collect data, so this is a I think this is actually listed as a as a big challenge, one of the big challenges for nsf cut the data revolution, and this is squarely in the middle of that kind of challenge.

326
01:01:54.480 --> 01:01:54.840
Thank you.

327
01:01:56.130 --> 01:02:06.660
Manish Parashar: here's a question from haunts krim how are the edge units instances powered how critical of power concerns in the device architecture.

328
01:02:08.250 --> 01:02:10.710
Peter Beckman: Really really good question and.

329
01:02:11.910 --> 01:02:23.820
Peter Beckman: So we have been very fortunate that in the places we're currently deploying there is line power and city of Chicago, for example, gives us a 10 amp circuit.

330
01:02:24.570 --> 01:02:33.570
Peter Beckman: To us, but we know that that's an impediment, so we are, we have a very clear power budget.

331
01:02:34.440 --> 01:02:40.530
Peter Beckman: The indoor nodes those sage blades that going instrument huts when we talk to the.

332
01:02:40.920 --> 01:02:47.790
Peter Beckman: People who haven't from in huts they tell us well, we have 100 Watts, you know, or we have less or we have 200 but you can't.

333
01:02:48.060 --> 01:03:00.570
Peter Beckman: And the temperature gets to be this temperature and then we're going to cut things off right so power is is critical in the in the independent nodes that sit outside you know 5070 Watts is what we've designed for.

334
01:03:01.020 --> 01:03:11.010
Peter Beckman: But one of the goals that we have we've been asked by our collaborators in Oregon to design a version that would use 48 volts right or 24 volts.

335
01:03:12.060 --> 01:03:21.270
Peter Beckman: me, and that would allow us to run off of places that are even more remote that are solar powered and so maybe even for the Hawaii deployment.

336
01:03:21.960 --> 01:03:34.470
Peter Beckman: We may need to look at what it would take to adjust and our plan and see if there's a way to be powered off of solar, so we can get a 3040 Watt package deployed.

337
01:03:36.870 --> 01:03:50.970
Manish Parashar: Thank you, we have another question from Lucia Zhang given the scale and multi tenant nature of the system wonder what kind of security solutions have been built into the system.

338
01:03:51.660 --> 01:04:07.560
Peter Beckman: yeah thanks um another great question, so there are two parts to that i'll answer the first part i'll just repeat from a little bit explain a bit more what I said earlier, is because these nodes originally are going places, you know very remote.

339
01:04:09.000 --> 01:04:19.950
Peter Beckman: Even whether it's in an instrument hot or it's or it's outside they need sort of extreme reliability, so in the nodes that are designed for outside rush ash and the team have.

340
01:04:20.340 --> 01:04:36.240
Peter Beckman: You know, multiple booting that can happen if it can't boot this image it falls back and wait some period of time and boots another image and so there's high resilience, but on the cyber security side they only phone home right so so we're very.

341
01:04:36.990 --> 01:04:46.590
Peter Beckman: You know you can't ever be too confident with cybersecurity, but we feel pretty good about the fact that it's hard to penetrate because the place of penetration is the.

342
01:04:47.100 --> 01:04:55.140
Peter Beckman: The servers in the cloud it at northwestern and you would have to go through the northwestern cyber security issues, because these things are only.

343
01:04:55.500 --> 01:05:03.360
Peter Beckman: phoning home, and then they opened a bi directional communication, but it's only phoning home, but from the cyber security of the.

344
01:05:04.290 --> 01:05:14.010
Peter Beckman: data and the running code that's a harder challenge you might know that docker and other containerization technologies.

345
01:05:14.700 --> 01:05:29.010
Peter Beckman: suffer from weaknesses in cyber security and so multi tenancy means for us that we have two tenants, who are willing to share now they get separate docker instances and and they're separate but.

346
01:05:29.520 --> 01:05:36.000
Peter Beckman: You know, they could interfere with each other, so we assume that they're not adversarial in sharing the node.

347
01:05:36.750 --> 01:05:48.870
Peter Beckman: Because you know you are sharing a scientific instrument and that's not a I don't think that's outlandish to to start with that position, which is that the people using the instruments are not adversarial.

348
01:05:49.770 --> 01:06:01.170
Peter Beckman: Who are given accounts right to use the use the edge so that's our current position, but we believe there's a lot of R amp D space to to evolve that and we've talked to scientists at northwestern and at.

349

01:06:02.310 --> 01:06:04.380
Peter Beckman: Pittsburgh and other places that look at.

350
01:06:05.700 --> 01:06:19.860
Peter Beckman: Technology that would use the encryption capabilities of the modern chips and the safe areas of the modern chips to allow for components to be locked down so that no one could look inside your container.

351
01:06:22.170 --> 01:06:36.150
Manish Parashar: Thank Thank you pete we have another question from Donald walsch all the data taken shared so that people can use it for Meta analysis benchmarking do the projects also share code.

352
01:06:37.320 --> 01:06:46.890
Peter Beckman: Yes, both great question, so the first question is, if you go to our github site for sage and waggle.

353
01:06:47.400 --> 01:06:58.800
Peter Beckman: All the code is open source all of it is open source, if you go to portal dot sage continuum.org you can see all 12 of those Ai Apps and you can download them.

354
01:06:59.490 --> 01:07:11.820
Peter Beckman: It takes you right to their github page, so one of the tenants of this kind of cyber infrastructure is you're building systems so that other people can can leverage them and build on top of them.

355
01:07:12.210 --> 01:07:28.800
Peter Beckman: And so, that means they have to be open and we've we've started all of our conversations with team saying we have to have open code now on the data side, all the data is shared immediately to the cloud it's open, we have two exceptions, however.

356
01:07:30.000 --> 01:07:42.090
Peter Beckman: So one is there are times, where we collect training data which is data that we're going to later use for Ai algorithms to build models and that data should not be.

357
01:07:42.570 --> 01:07:51.990

Peter Beckman: Put posted on the Internet examples are from the city of Chicago we have an agreement with the city of Chicago that if we occasionally take a picture of.

358
01:07:52.290 --> 01:08:03.450
Peter Beckman: pedestrian movement, so that we can then improve our algorithms in calculating pedestrian traffic that's fine, we can share that with other researchers, but we're not going to post that to the world.

359
01:08:03.960 --> 01:08:10.410
Peter Beckman: So those kind of training data is our policy protective and the second case is with the.

360
01:08:11.310 --> 01:08:20.700
Peter Beckman: native tribes the sovereign tribes, and there are some data that might turn out to be culturally sensitive and so we're.

361
01:08:20.940 --> 01:08:35.760
Peter Beckman: cautious and we work with them very closely to make sure that we can understand how that data can be stored where it might be stored, it might not be stored on our servers and might be stored on their servers so we're trying to work that out with them.

362
01:08:38.130 --> 01:08:49.200
Manish Parashar: Thank you can bang ask do you foresee application of mobile sensors instruments in space cetera what new issues would arise.

363
01:08:50.850 --> 01:09:07.650
Peter Beckman: yeah that's a that's a fun one because I don't know how much of my office, you can see, but you know we like to take her, and so you know whether the drones or or other mobile robots we were very interested in.

364
01:09:08.700 --> 01:09:14.820
Peter Beckman: In what what kind of processing, we could do at the edge and there already have been groups that have done things like.

365
01:09:15.360 --> 01:09:29.430

Peter Beckman: built a drone that has a nvidia gpu and that nvidia gpu is processing looking for power lines, so the idea would be you would launch the drone and have it follow the power lines.

366
01:09:29.730 --> 01:09:36.210
Peter Beckman: Especially after a tornado or some event and just map them what's happened to that right get video feed of that.

367
01:09:36.540 --> 01:09:46.140
Peter Beckman: That doesn't need a human operator to you know pilot so that Ai at the edge, we believe sage could provide a very interesting.

368
01:09:46.800 --> 01:09:57.390
Peter Beckman: space for that, where you again, you could write your docker container ship it off to the drone the fleet of drones and change it up okay we're now looking for a particular.

369
01:09:57.660 --> 01:10:06.270
Peter Beckman: Rescue truck that we haven't been able to contact you know, this is the image of the rest of the truck so if any of the drones sees this rescue truck please land.

370
01:10:06.540 --> 01:10:24.420
Peter Beckman: and set up communications back right there are all sorts of scenarios where mobile processing like that might be very interesting and i'll mention one other thing that we've are starting to explore again in the hobby sense is mobile and satellite communication so.

371
01:10:25.800 --> 01:10:39.000
Peter Beckman: there's there are several companies swarm dot space, which was just bought by spacex have have Pico sats and they you can you know burp up small packets of data.

372
01:10:39.570 --> 01:10:49.080
Peter Beckman: So that picture here of Aaron and bill might be later in the middle of Alaska or the Amazon, and all the processing.

373
01:10:49.350 --> 01:10:58.350
Peter Beckman: Your did you see this species what's happening on the ground what's happening with flooding could be reduced and analyzed into just a few packets that you could.

374
01:10:58.740 --> 01:11:05.040
Peter Beckman: You know spit up to the to the satellite, so we have one of those tests units that's on my kitchen counter at home and.

375
01:11:05.580 --> 01:11:15.060
Peter Beckman: i'm hoping to get some time you know hacking on that, and in fact satellites have been doing a kind of edge processing to because they can't possibly pull down all of the data.

376
01:11:15.390 --> 01:11:24.570
Peter Beckman: And so they look for particular events and process it, but again power is a big problem in space so it's ultra ultra low power process of their.

377
01:11:26.490 --> 01:11:27.810
Manish Parashar: Okay, thank you.

378
01:11:29.100 --> 01:11:37.800
Manish Parashar: We have an anonymous question that starts by saying great presentation appreciate the focus on complex systems in the outside environment.

379
01:11:38.400 --> 01:11:52.470
Manish Parashar: Do you see similar opportunities or challenges for edge cloud computing in network manufacturing iot environments, perhaps across different companies their privacy is a significant concern.

380
01:11:54.690 --> 01:11:59.880
Peter Beckman: Yes, I would say that privacy comes up very frequently.

381
01:12:00.750 --> 01:12:06.210
Peter Beckman: And it's a good thing, I mean I really think it's a great way for us to understand.

382
01:12:06.510 --> 01:12:13.440
Peter Beckman: What data is being collected and why and edge computing gives us a way to get our hands wrapped around that one is.

383
01:12:13.680 --> 01:12:29.730
Peter Beckman: All the up in our case, all the algorithms are public and all the processing public, we can see the codes, we know what is being processed at the edge, but you know I would love to see that kind of rigor applied everywhere, you know i've been to.

384
01:12:31.290 --> 01:12:43.980
Peter Beckman: You know the new kinds of stores, where they look at video and how you move and understand whether you purchase something or intend to purchase something that in manufacturing it's happening all the time.

385
01:12:44.670 --> 01:12:59.310
Peter Beckman: This issue is one where there's our end that's very key and very important here, and I know that nsf has projects in the privacy space, and this is a insecurity, this is a.

386
01:12:59.790 --> 01:13:11.580
Peter Beckman: Very ripe computer science problem to address and the kinds of processing that we can do on the very smallest devices and iot devices, is also an area of interest.

387
01:13:11.940 --> 01:13:20.160
Peter Beckman: Whether it's just anomaly detection so devices, they can't run a full Linux docker container but maybe there's a little snippet of code just.

388
01:13:20.970 --> 01:13:34.470
Peter Beckman: A bit that can run with a new kind of algorithm that does anomaly detection that looks for a certain vibration, or a certain sense that something is amiss so there's a lot of a lot of room for innovation there.

389
01:13:36.540 --> 01:13:37.380
Manish Parashar: We have one.

390
01:13:38.610 --> 01:13:45.840
Manish Parashar: More question from the initiation could you share your thoughts on storage support in the long run.

391

01:13:46.740 --> 01:13:48.240
Manish Parashar: Supposedly, the amount of.

392
01:13:49.050 --> 01:13:54.180
Manish Parashar: syncing data will grow exponentially who or how would.

393
01:13:56.190 --> 01:14:01.500
Manish Parashar: How the long term data storage costs covered, what about the agility of the data.

394
01:14:02.460 --> 01:14:22.560
Peter Beckman: yeah Thank you very much, and I appreciate that question because I forgot to mention it my talk so I feel really bad now um so one of our partners who were where we're storing the data, in addition to northwestern and San Diego has been osm open storage network Alex delay.

395
01:14:23.760 --> 01:14:39.480
Peter Beckman: These the storage pods that provide an object store and so right now, all of the sage data is at the object store in San Diego and, in this case and and the team there, allow us and give us.

396
01:14:39.900 --> 01:14:48.540
Peter Beckman: terabytes and terabytes of space, thankfully, the relationship with them has been very good we're very thankful to the sponsors of that which is both nsf and.

397
01:14:49.890 --> 01:14:50.940
Peter Beckman: Private donors.

398
01:14:51.990 --> 01:15:01.560
Peter Beckman: And so right now we've asked for space there and expanded that space and the team has said, well you've asked for a little bit more, you might have.

399
01:15:01.980 --> 01:15:14.070
Peter Beckman: eventually have to start to ask for a lot more that that one high rez burn that we did at the prairie generated 250 300 gigabytes of data, and that was just a couple paper.

400
01:15:14.400 --> 01:15:18.990
Peter Beckman: So if you really tried to store it all, which you can't that's why it's computing necessary.

401
01:15:19.410 --> 01:15:27.690
Peter Beckman: But even so it's going to be a lot of data, and so the lesson is a great model we love that model where they just allow us to.

402
01:15:28.170 --> 01:15:41.400
Peter Beckman: Make objects and store them in their object store, we also assume that for longevity, we need to store things for five and 10 years and so that's part actually part of nsf.

403
01:15:41.760 --> 01:15:54.690
Peter Beckman: Planning always is what's that what's going to happen with that data but you're right, it does grow exponentially it's a time series right so it's just going to keep growing as long as the instrument is there and, especially, as you add high resolution sensors.

404
01:15:56.850 --> 01:16:17.580
Manish Parashar: Thank you, we have one final comment from Michael Boucher and his comment really reflects all our thinking and it says, a big thank you to pete for such a incredible talk and for answering the questions so with that, let me give a unless there any other questions.

405
01:16:17.670 --> 01:16:19.710
Peter Beckman: i'll share a bonus slide.

406
01:16:20.850 --> 01:16:23.910
Peter Beckman: So just for fun i'll share a slide.

407
01:16:25.920 --> 01:16:27.960
Peter Beckman: See here so.

408
01:16:29.130 --> 01:16:31.320
Peter Beckman: When after we had several nodes up.

409
01:16:32.550 --> 01:16:44.280
Peter Beckman: We were talking about the tragic eruption of the of the volcano the under water volcano and and our collaborators and Oregon said.

410
01:16:45.870 --> 01:16:56.160
Peter Beckman: Did you guys detect anything and and we you know we have sensors that are pretty simple sensors we didn't really think about it much and we said, well, we should go look.

411
01:16:56.520 --> 01:17:07.440
Peter Beckman: So Raj ash went off and looked at some of the sensors we had one in Oklahoma Texas in the one in utah was not quite up and running, yet we had one in Oregon.

412
01:17:08.460 --> 01:17:12.570
Peter Beckman: And Sure enough, even just with a simple $3.

413
01:17:13.890 --> 01:17:28.170
Peter Beckman: Air pressure sensor just to you know pressure humidity temperature kind of sensor, we can see that wave the pressure wave as it rotated and bounced around the planet, a couple times and this really.

414
01:17:28.770 --> 01:17:47.820
Peter Beckman: goes to the point I mentioned about anomaly detection after this happened we realized crud, we need a set of algorithms that would have seen this discontinuity and running that at the edge 24 seven so when we see a pressure wave it immediately alerts us right, this is the kind of.

415
01:17:49.050 --> 01:17:52.440
Peter Beckman: Multi messenger science that that physicists do right.

416
01:17:52.740 --> 01:18:02.730
Peter Beckman: Where you notice an anomaly, and you can report to other instruments, so we want to add that anomaly detection that's a new area that we would love to do, but it also shows the power of.

417
01:18:03.000 --> 01:18:11.010

Peter Beckman: This kind of gathering and what we could do with processing at the edge and so with that i'll i'll wrap up.

418
01:18:13.230 --> 01:18:32.460
Manish Parashar: Thank you again that this is fascinating really, thank you for taking the time we appreciate your talk and thanks everybody for attending and I just want to remind our nsf colleagues that there is a QA session later in the afternoon with pete Thank you.

419
01:18:33.300 --> 01:18:36.480
Peter Beckman: Thank you very much thank you bye.