

Supermicro A+ Roadmap

0:00

Broadcast is now starting, all attendees are in listen only mode.

0:05

Well, good morning everybody. Welcome to Day four of ASI's 2020 Winter Technology Summit.

0:14

I want to thank you all for joining us again today for Friday and for making it through our week of seminar sessions.

0:24

Today we're going to be closing. Now we have our cleanup or our cleanup hitter today.

0:29

We have Supermicro and we have John Park from Supermicro who's going to be talking about those Supermicro AMD product lineup that they have and we'll get started with his presentation in just a second. A few quick housekeeping items that I wanted to update you guys on. First, at the end of today's seminar.

0:52

I just want to ask everybody to kinda hang with us for a few minutes at the end. I want to kind of have a few seconds to wrap things up and talk about the week in general. So we'll do a quick wrap at the end.

1:04

But just to remind everyone that supermicro, for their seminar session today, they're going to be giving away a pair of beats over the ear headphones, And we're also going to be raffling off a bunch of \$50 American Express Gift Cards.

1:23

We've got, you know, 10 of those that will be given away, along with the headphones, for lots of cool stuff, To raffle at the end, for your Thank you, for your guys, for participating with us today. So we'll be doing that.

1:35

From yesterday, we had the Samsung seminar session, and their presentation I talked about their FSB product.

1:43

And their Raffle Prize was the 50 inch LED Smart TV, which we said, we announced the winner of that this morning before we started supermicro. So I'm going to go ahead and do that right now.

1:57

I want to congratulate and ask everybody to put your virtual hands together for Carl Hogue from Wisconsin. So, Carl, congratulations!

2:07

You won the Samsung 50 inch LED Smart Display. So, way to go, Carl.

2:15

Also, last thing, the grand prize that we're giving away is a Intel gaming notebook.

2:22

So, for those of you that have joined us through all four of the sessions, you're going to be entered into the raffle for a chance to win that notebook. So, lots of cool stuff.

2:32

There'll be a lot of e-mails from me coming out after the event is over, so I'll apologize in advance for bombarding you guys, but I got a lot of information, I need to get out to you after everything is finished.

2:45

So just do expect a lot of follow up e-mails from me after the event, But with that all said, I'm going to turn it over to John here in just a second, But I want to remind everybody that for questions, you can submit your questions by typing it in. There's a question icon on your control panel. You can just click that and open the window and go ahead and type in your question.

3:11

And then at the end, we'll make sure that we ask the team from supermicro your questions. So go ahead and type those in, and we'll make sure that we get everything answered at the end.

3:21

So, now, what that's that, Qian, it's all yours. Take it away.

3:28

Thank you. Thank you so much. Hey, hi. My name is John ..., System ... for micro, mainly in buckets of the aim, the base product lines.

3:37

So today, I have a very broad portfolio of AMD based servers that is, wait what we have today from the simple micro.

3:51

So we have a lot of slides today mainly for the basic aspects of servers, and each server is a unique, too different applications and whatnot. So if you have any questions, and I'm assuming that you guys will have a lot of questions and please, write those down, and then we can be able to answer that as we go, or we can wait until the end of the presentation, OK?

4:16

So, brief history of a AMD product, or a media projects or roadmaps. So, from the 2017, the first and April was launched until today.

4:28

We are seeing a Milan coming up anytime soon. But that don't let that eight, will you use 20 20? But it's actually unknown.

4:37

So, other than those, this is the, basically, a very brief product line of what the AMD has been doing.

4:45

Give you a quick update on the, What is the difference between them, and they put them the wrong.

4:50

On April was a first generation ...

4:53

processor that was newly developed from the ground up.

4:59

As you all know, that the Intel was a predominantly dominated the market CPU markets from the 2000, I will just say, middle of a 2000.

5:09

Or so until still today, but Maple was a challenged by the Intel. And there was a lot of good features that was in that particular Naples TPUs.

5:21

Kicked off kind of away from the Naples, The Rome wasn't that. So it has a lot of differences between the maples and alum.

5:27

So we have some challenges to educate our customers and delve into these.

5:32

So, the applications that wasn't usually ran for the Intel based service in the past, he wasn't really working so well on the maple side because during the absence of the AMD bases database, all the applications was a weekend based on the Intel eight.

5:49

So, when the customer pool trying to adapt to the AMD, they had a struggles because it wasn't money. But ways of post the.

5:56

right, All of those was ironed out from the year 2017, and 2017, 2018, Now, that the wrong was a, basically fully mature systems.

6:09

We were happy to announce with our H 11 based servers in each default ASAP.

6:15

So what it means is that ... 11 days of the cerebral ...

6:18

while the forest launched with a nap over.

6:22

So, I love the generations of so when maple was first launched H 11 server was developed with the categories like a big please WIOA Superstar Rachel chai empowered.

6:37

If you are not familiar with the our naming conventions Daniel just go along with me. I'll have a lot of presentation and that can kind of hint in a way that you can be able to understand what those what those are and the age of 12 days etcetera is geared towards for the Rome and Milan. It. It's a socket compatible so we can be able to dropping it with the iOS upgrade.

6:59

And H 12 H 12 servers.

7:02

The key key key things to point out is the ... for, where I asked the H 11 is just the remaining ... three.

7:14

Even with the wrong dropping, you can't be able to get the 64 cores 3200 MHz and a memory channels from the H 11 servers and H 12 servers.

7:25

Only difference is that the because of the PCH E is hard wired into the motherboard.

7:32

So, a PCI Chancery still remain on H 11 with a long dropping. However, on the 8 to 12 style, you get the full benefits of what the ...

7:42

post, OK? So, brief highlight of what we currently have today.

7:48

So, I can be able to categorize into the three different sectors. Rack mounted a multi-million, a GPU optimized servers.

7:57

So what I mean by direct noun is that basically, 19 inch.

8:01

You know, you put U 1 U 2 U when you are to your form factor, it could be a single psaki, a dual zakat, one day per channel design or 20 per channel.

8:12

There's a various of The combination is on the status ad, audio, you that to MBA me. Along with the onboard M dot two. And then the redundant power supplies now on that. So all those good aspects of the server or we have in the roadmap as we speak.

8:29

But right now, as you can see that the Ultra system, we have a 1 U 2 U 1 factors, and then the WI Oh, we have a one day of a channel design lawyer or to the Channel Design and then the multi multi site.

8:45

We have a 2 U 4 node with a single socket, or two, you will note that dual socket, and then we have a leap on factor with an eight you want factor with E at 20. 20 systems, you know, a phone factor and then we also have a GPU optimized airbrushed like PCIE version of the GPUs.

9:04

Or the MVD rests on an adult, that type of GPUs.

9:08

So having said all those, what we have coming, I'm not going to spend a lot of time on this particular product, because it's not readily available today.

9:18

But we have a spinoff product lines of light on Ultra WIOA and a path to N or multi node and that you optimize and the meat power servers.

9:27

So, I'll start off with, up with the rack mount servers.

9:32

So, um, far as we like to emphasize a ..., so, what even Alceste has been, is that the, basically, what the AMD CPUs as pieces of what we can be able to do. We can't be able to emphasize all those key features into our servers.

9:49

So, I'm not gonna go in bullet by bullet, but the key highlight of this particular server is that the 1 U 2 U. one factor of 20 per channel design with the two socket covers up to the eight terabytes of memory is of 220.

10:03

I mean, sorry, eight terabytes of memory would be highest the kickoff capacity of the teams today are available. We support degree here.

10:12

Good thing about the to your phone factor is that the because they have a lot of a PCIE coming out at the back, we will be able to support the double width GPUs in here.

10:23

So if your customers are our customers or anybody else who has an interest in double with GPUs, maximum up to two, we can be able to provide that with the supporting of the highest the PDE CPUs, which is at 280 watt CPUs.

10:38

We can be able to provide that in here.

10:40

one thing that I also wanted to mention is that this particular product has Dina flying off the shelf ever since he was born. Because it has a capability of doing A and B and me as well.

10:53

So, a lot of the customer who just wants us data, that'd be happy with this particular product line. At the same time, there are other customers always interesting to phone. Factor of an MBA me and it will be able to utilize this particular one, but not all the back lanes.

11:10

Slotkin do support the MBMD.

11:12

So, that is the reason why we came up with all direct attached NBA meets up version of the Ultra.

11:19

Previous generations of our age 11, when you form factor. We won't be able to solve for up to 10 youth acting

11:28

The highlight on this one is that you got to ... support on one.

11:32

You on the aged 12 version is the 12 drug base.

11:36

And then be to your form factor. It's all direct attached and pay me.

11:41

So, that's the main reason why the Treaty of Compact that can only support one, the PCI 4 by 16 coming out of the fact that because individual you, you, doctor NVME takes up to four PCIE lanes. That is the main key features. So that's not a typo. And then when you form factor, we have a tunnel, PCIE.

12:08

If you count those, we have a four PCI coming out of the back, all the Power supply, side redundant, when you to your form factor is still using the similar motherboard as previous one and this particular one is a good for people who will be using AI Ops, i-hop generated, or somehow somebody with that I have already memory intensive applications.

12:35

So, we have been seeing a lot of customers who are requesting this because of the, uh, sufficient power supplies with the memory to core ratio, plus the Indian. So, we've been seeing a lot of requests coming out of that area as well.

12:52

And then V A one U factor of the WIOA in one day for Channel design, is that the, a lot of a customer who was lovely, looking for a cost effective, a version of the data center environment.

13:06

Where they can able to utilize up to 64 cores of the CPUs, and then be able to accommodate up to, took terabytes of memory, and a plus B, three PCI events coming out at the back at the same time. This particular one can be able to support up to two nvidia, cheap for GPUs.

13:27

So that is the one of the key features, the reasons why we have a lot of a customer who used to be stuck at the, to the power, sorry, two socket version of the system.

13:38

They loved this particular product line, because it does not picked up a whole lot of a power. And at same time, if the when you plug factor and the app, and it can be able to provide you with a three PCIE coming out of the back. So a customer who was interested in in a high bandwidth and networking, like 100 G, 2, 1 G, networking, car, MP four, and then add a T four. So they deem that really, really happy with this particular product. So, this was a full surprise, surprise, when we first launched this product, that we had high hopes, but then exceeded our expectations because the data center people literally fell in love this product from the get go.

14:18

one new form factor would be a 43.5 an inch dry bay can be converted to a for you, doctor

14:27

And then the 2.5 is dry bay. Can't convert to You'd have to NVME. So, that is the key, take away from this particular one.

14:37

And then we have a up the annual a little bit, And this particular product is the same that we got your product line, but it's a two input channel design without all direct attached NVME.

14:48

So the reason why this was born is because we had a good run with the previous generations of gays, a lot of customer who wasn't looking for a memory intensive combinations and would be enough for count plus the NVME. So they really love to continue to see these in a PCI slash and four. So this is the reason why that we were able to kick things off with the could trend.

15:18

So a lot of a customer who is still looking for a 24 ... A, and this has to be a key aspect of their application.

15:26

At the same time, we can be able to provide, we'd be one PCI ... as well. So I'll get to the same concept as our ..., all the PSAP's were saturated. Supporting the Albright attached NVME.

15:43

So that is the reason why we only have a one PCI coming out of the back.

15:49

So I'm going to pause here for a little bit, and I want to ask your question, or I want to take a question because I wanted to understand the if you guys understand what I meant by all direct attached, MBA me, and that's the only reason why we can't be able to provide a one PCIT lane.

16:12

So, can are you asking if there are any questions, or Yes, I am.

16:18

Yes, I am asking questions. If that's not a confusion because some customers that I have faced in the past whenever I ask that whenever I tell them the questions.

16:27

I mean, whenever I tell them the answers or I make a statement that the 24 direct attached NVME. Why is that Why is that that we can only have a lumped, PCIE?

16:39

So, that has been a some customers were asking questions in the past.

16:45

So, I was wondering if any of the audience have the same question?

16:49

Well, they haven't asked that specific question.

16:52

We do have a question asking about the 24 Bay, two R U and V M E server wanting to know if that supports hardware raid for you to NVME, not software raid but specifically raid five, raid six for the hardware?

17:12

Great, so, sorry that I didn't mention anything about the right configuration.

17:16

Very good question.

17:17

So, as of today, and all you know, the AMD does not have the control of the CPU.

17:24

So, what is going on is that we have to utilize the tri mode rate controllers.

17:31

But unfortunately, as of today, we cannot do a raid at using a tri loan rate controller.

17:41

As of today, the status asked was fine with that particular timeout.

17:46

But, you, doctor ..., it doesn't seem to be working so well at the moment.

17:52

So we are waiting for the next generation of the controller.

17:57

I don't have an answer or the ETA when that's an idea available and then one that's gonna be invalidated, but that is a known issue, and then that we are addressing that as we speak.

18:11

Then the Ultra system, that's the ..., the next one is that cloud DC.

18:17

So it's a very similar to what I've shown on the WI.

18:21

Oh, one of the biggest takeaways from this particular platform is that it has the company called a IOM.

18:30

What the AIM is, the Advanced IO module, it's the OSTP three Compliance.

18:36

So a lot of a customer who wants to upgrade their networking devices, majority of the times that the networking devices were on board devices.

18:48

So, if you will, you need to swap out the motherboard in order to upgrade the onboard network cards.

18:57

In this particular design, that a.i.y. means that piece aggregating design, where you can be able to pick that portion out alone.

19:05

And then you can be able to upgrade to the next speed of the net or next to a card that you would like to install. Example, least like a four port, RJ 45 point factor next door.

19:17

If you want to move it up to 25 G, then you can be able to do that just by I'm populating via a slot. And then that populate the new cards that you, of your desire.

19:29

At the same time, when you contact that can be able to provide you with the two PCIE by 16.

19:35

So E addition to the neat cards or ..., you can be able to support the two extra PCIT devices into cyber. It's a very similar to what you have seen on the WIOA on it. The one of the biggest difference is that the WIOA versus the cloud.

19:53

you see that I just data version with the ...

19:57

Channel Design, and it also supports the ASP 2600 So, that the R T root of trust is enabled on this particular product line. That is the key takeaway from this particular product. And we have a lot of hope on this one, because we have been asked the questions on the WIOA, our applications, and some meeting a more than just the three PCIE. We are using it two spots for the neat cards and then we are requests to have on other spots for the our other devices. And this is a perfect marriage for that combination. So we are looking forward to the particular one.

20:38

Now that we have to you form factor of less similar design, the reason why this is the very popular product is because we've been asked that a lot of times that I'm looking for a single socket, 20 per channel design, that can be able to support 1 or 2 GPUs.

20:57

So they are talking about the double loop GPUs or the T bars.

21:01

So we were able to provide, we'd be a tooth double with GPUs or the four G force in this particular product. It's the default as a data.

21:12

And then this one also supports the root of trust.

21:15

And then this particular product can be able to support up to the four terabytes of a memory.

21:23

We are looking for a good balance between the GPU and the CPU cores, and then the memory on this particular product wire.

21:34

Next one will be the To your phone factor with a single socket.

21:43

We don't really have a category for this. And we would like to call these a mainstream is what we used to be called.

21:51

So this particular product is a brand new, we just kick this thing off.

21:55

And one good thing about this product is that it had the event of a one day per channel design fully populate with the established PCIE by 16 and by a combination of which.

22:08

And then B also has a on what? 45.5 coming out of the Mac.

22:14

So, we have a lot of customers who has been asking about the, I need to have a lot of a neat cards in our system.

22:21

So, this is one of the main reasons why this particular product is born, and a lot of the customer who's been teaching for a mainstream product line will be happy to see a two year bond factor, a single socket, one different channel design on this particular line.

22:38

But that's the essence of the rack mount server.

22:42

We're going to be moving onto the multi node, in this particular case, so multi node. I want to give a little bit of a hint on a wide. The more I know it has been so popular from. The get go.

22:55

Is because it's thought that everybody knows you, regardless to Intel, all of the A and B, the end product, were the people who was in the data center, where they can be able to, you know, have as many nodes as possible in a smallest form, are smaller.

23:13

Rack footprints, people's, again, looking for something like these.

23:16

And it's been known HPC markets as being very popular in the T four nodes because the HPS and Market does not require a whole lot of the drives, but they would like to have a minimum PCIE.

23:28

And then at the same time, they like to support a balance of the memory to the CPUs.

23:37

So, because of that reason we excel well in our previous generations of the between product line. And as you can see on this particular product you have two different feature sets on this. one is the all data version.

23:51

Omega one is the Hybrid President supporting the four NVME and tools SATA.

23:58

So, impose your Y that shows the enclosure packed planes and power is the power supply and the Individual Node Act as if different, just the individual server.

24:10

The good thing about this one is that the deep dive with the S I O, am not the ILM differences.

24:17

between the S I O M and A I miss that S IOM's supermicro IO module.

24:24

It's a proprietary cars at the right, a ILM.

24:28

I'd be a disability not, or CP compliant.

24:32

So it's a proprietary design is aggregated design I mentioned before in case that customer wants to upgrade their network in devices without the replacing the motherboard.

24:44

They can simply swap the IOM and then they get the desired reports that they want.

24:50

So because this is so dense, key thing that I want to point out is that the not all CPU can be supported in the air cooled.

25:00

Um, some of the, some of the CPUs that we have seen before, like seven app series from the Rome series and then the 7 age 12 from the wrong.

25:14

Those have a higher GDP, 240 and then 280 watts.

25:18

So we have some limitations and what can support, eat and aware?

25:23

So some customers don't care for those, because they don't always looking for the highest TDP CPUs. But we've been pretty happy with what we have seen.

25:34

And this has been a one of our flagship Taiwan's and from the Intel and AMD side.

25:40

And we have been expected any meeting that expectations are with this particular PowerPoint.

25:49

Next one is the sister version of the Big 20, it's called a ... Pro. A lot of a customer who's in the HPC market Older Data Center has been asking us, hey, we don't need to stop. And we are happy with the 64 cores and we're happy with them making them be in pounds. And what we need is something very simple, cost effective version yet, it's a multi node.

26:12

So there are a lot of a customer goes into a single socket, one day, per channel design, does not require a hard drive or whatnot, hasn't been very happy with the product because it has an A SI O.

26:25

And just like what just their version has same enclosure shares via back plane and the power supply.

26:32

And a customer who's looking for a highest GDP is, be able to do that. So, in despite the one, at the same time, the Sony is equipped with a four on board and up to.

26:44

So the customer, who are looking for a very small footprint of the capacity at the hard drives, and this has been the platform for them.

26:54

So we're happy with this particular one because of a lot of, like I mentioned, a lot of the HPC market and data center people who doesn't want to spend too much money on a dual socket. This has been the product for them.

27:07

The bigger version of that is the Blake, so we have a two different blades.

27:13

And one is the non GPU supporting, and in this line, you see, this was a massive chief you're supporting.

27:20

And then the next slide that I'm already show, either GPS point, key takeaway, from this particular one, is the eight you pack with the 20 nodes individual, not act as just the regular server, and then he has the management that can be able to control as the chassis.

27:38

So we have a direct mail management tool that's separate from the PMI and then the week with the one, the PCI E, gen 4 by 16, and then, has a note.

27:52

And it has a built-in fans coming from the back, which is the hot swappable, and has a built in switch inside, and then also the ..., any cars, that's within the node itself.

28:06

It has a 25 G, or the 100 or 200 G capabilities on this particular one.

28:14

Next one is the GPU supporting.

28:16

So, it's, it's a little bit different from what you have seen before. What the key takeaway from this one is that we have been validating, and the base of GPUs, like ...

28:27

100 or the nvidia side. We are validating with a 100 and then that we are doing with an 8 40, and then the next generation of the GPU that's in line with an arrow and V.d.s. side.

28:40

So very similar concept of what you have seen in the previous slide.

28:44

And so this is the key takeaway from this particular one is that the GP Assemblies.

28:51

Next, couple of slides are the fact claims.

28:55

A lot of Accustom, I did ask me A, what is that?

28:59

key differences would be thrown by Factoring and a Tween Pro or the big point because it looks like it's the double the between or the Tween Pro.

29:10

So, fact when we have a two different flavors as a high level, one is the front row.

29:18

And then the other one is that we are IO I want to emphasize. What would it be here on the bad tween?

29:25

Is the individual node in this particular slide that you're seeing is a serviceable from the front. And in all, the iOS Index.

29:34

Only thing that you need the service is the fans and the power supply.

29:39

And this particular product lines back to impound line, the four nodes that you see on the left side, and the borno you see on the right side mechanically separated.

29:51

So, on other key differences between the between between Pro versus ... is that the eliminates the single point of failure when it comes to a hard drive back plane. In case of hard drive, that plane needs to be surveys.

30:08

The node E cell contains the path length, so you're not losing the entire neighbor neighboring nodes.

30:17

So I like the big points that I've had tween.

30:19

I'm sorry, Bitcoins that are too pro.

30:21

In order for you to ..., you have to power down entire chassis, whereas that Tween, you can only power down the problematic note If that exists, that's the honor, the key, takeaway from it.

30:36

And then, the, as you can see from the rear view, the two power supplies on the left and the two power supplies in the I was a kind of middle leash left to power supply, Hub plus the math nodes, and then the right side. And those are supported by the adjacent to that particular the past time.

30:57

So, it's a mechanically, sorry, electrically separate the in the middle. That is the one key takeaway that I wanted to point out.

31:04

This particular model, however, is the 24 channel design single socket, supports up to 280 watt.

31:13

See a TCP CPU.

31:14

So that's the main main key, takeaway from this one next to each E, we're IO. So what I mean by that is the, all the airports are in the rear, but still can be able to service from the front.

31:28

As I mentioned before, the, all the bad claim is that within the individual node.

31:34

So that's a lot of a customer who, you know, into just the solving the problematic and serving the problematic cyber only. They didn't really loving this particular product line because of that reason.

31:49

Um, next one is the two.

31:53

You kinda have a form factor, you know, at that point, a lot of our customers, who are using the ease of comparing with a Hadoop, a lot of a customer that even the Hadoop area, they have some kind of had a some kind of a ratio between how many hard drives, hardcore.

32:11

So, this one, if you look at it, has a two more hydride base are coming out of the back. So you have a ...

32:18

twin, the twin, the back, so total of eight hard drives parse system. That was the key takeaway from this particular to you.

32:28

Sorry for you for nodes that are aware of a lot of custom locals and looking for a density begin looking at a particular one of the key takeaways.

32:39

That 8, 3.5 inch Friday, so that the capacity, harddrive date can be bigger than zero point five inch, OK, so that's all the, um, more pay, no side.

32:51

And then we will be jumping to the GPU.

32:56

So the GPU I want to spend a little bit more time on this particular product is that traditional way of designing the server from the Naples and our Rome is using a 128 PCIE lane.

33:10

So we were all told that regardless of a single socket or the dual socket from the Epic CPU, they have 128 PCI lanes.

33:21

That's because a single socket contains 128.

33:25

Dual socket is because out of that 128 PCIE lanes, they split that into a two.

33:32

So 64 parsed CPU easily used as a communication link between the two CPUs so remaining 64 lanes parse.

33:43

CPU equates to 2 1 to 28 ..., So, regardless of where you have a single socket or dual socket from AMD, you'd have up to 128.

33:56

However, this particular design, the communication link between the two CPUs on the 64 lane parse CPU, they group that into a 4 by 16.

34:08

So, the communication lane of a four X TMI, cross Global Memory Interface is the communications from book to CP.

34:19

So there were 4, 4, 4 X TMI in traditional design. However, this one piece of connected the one of the

34:29

So provide that to a PCIE support.

34:32

So that is the key takeaway from this particular design.

34:36

So, eight of the PCIE was directly designed to support the GPUs, that with ... boards, double with GPUs.

34:48

And then the remaining 32 GPU Site 32, a PCI lane, can be flexible enough to support a either the MBMD or the by 16 plus B by eight, or a true size 16 for the ALC card. So, without the NVME. So there are a lot of combinations you can be able to do with this particular product.

35:14

If you are not using the all eight attached to GPUs, and you can use those PCI spots or some other devices as well.

35:23

So this has just happened to be a one of our flagship right now, because the 160 PCIE lanes that's featured in here has been a really, really good features that for a lot of customers. So who ever faced in the past?

35:44

So because of the ... form factor, we don't have any problem with the supporting 280 watt GDP.

35:53

CPUs is a double with a channel design. And then beta has a option of the supporting and ..., as well.

36:03

So that a lot of our customers. So who wants to use a ... off the shelf? They can be able to do that using the product.

36:12

Current design on this one has a 24 dry base up in the front for the Which can be converted to an NVME.

36:21

And the remaining tri gate can be supported using a AOC cards like a 24 port, um, data controller or whatnot.

36:31

So you can be able to do that by utilizing all be a piece, the H harddrive slots and the font.

36:39

Keep in mind that this one is a limited to support up to four you, doctor NVME. That's because of the pack plane limitation. So we are looking into supporting up to 24 MBA means by swapping out the pipeline.

36:57

So that is we are in discussions to do that right now.

37:01

If we do that, then the 24 ... plus the sum of the Slack can be used for the GPU's older PCIE as well.

37:09

So we are looking for this particular product line.

37:12

That is kind of a spin off the support at the front applications in the near future.

37:19

Next one is the Redstone.

37:21

So rest phone is the F X M four days GPUs coming from the Nvidia.

37:30

So it's like a mother board looking with the embedded a GPU. So that shifting the four key things that you see on the left-hand side.

37:40

That particular shifting East Coast. Or the FSM for, base, the GPUs, and then the two other ... in the back side, that's for the CPU side.

37:51

Um, when this was introduced to us by Nvidia, they recommended us use it in a 3, 3, or four factors rather than have to you.

38:01

So about our design of the chassis, and the cooling the bands, and all the other combination, we will be able to accommodate that into a two year form factors, rather than a three you contact us.

38:13

So we have customers are lined up already to use this particular product because of the form factor has an advantage versus what the aim Nvidia has proposed.

38:26

And then that these velocities are all direct attached on the back end for ... GPU.

38:32

So there has been a customer who are looking for an upgrade there, ask them to, to SSM for it.

38:42

So we have a lot of customization who's been looking into this because of that reason.

38:47

Then the brother version of that is of Delta.

38:50

So it's a very similar to what you have seen them.

38:52

But the rest of outside, instead of having a four X M four version of the GPUs, this one is actually flip with eight SEMs for GPU. So, the key feature sets a very similar to the rest, don't.

39:08

This one has a great racial from the GP Ratios, from the need for the Alex ... design.

39:17

So people who have been looking into those applications has been using a lot of academia research and academia, and then even the HPC market, I can't name all of those customers, but has been already lined up to get a P of gets us systems that already, for that matter, key takeaway from that is the addition to all the other good features, we have up to M back to onboard.

39:42

So that is another key feature on that.

39:44

Lastly, we have an upcoming to you, two nodes that supports the three double with GPUs.

39:53

The reason why diesel wasn't born is because a lot of a customer who are seeking for a one-year form factor that supports two or more double with GPUs.

40:05

The reason why we made this scene to a kind of multi node is because we thought it wasn't more cost effective if we can be able to use a shared components.

40:16

So that is the reason why we come up with the 2 U 2 nodes rather than a pizza box version.

40:23

A lot of customers, or who's been looking for a one-year form factor with the single stock.

40:29

It doesn't matter that being pounds attendee wants to have a 1 or 2 GPUs. And they are, you know, in mind to get this particular one at the same time, because we are sharing the fans and enclosures and whatnot. So the cost is gonna be going down, just a little.

40:48

So, we have a highly anticipated in the HBC market for this one.

40:54

At the same time, we are looking for a people who was into the AI and Machine Learning and all those good areas that we are looking into, the support, those customers. And then, the one of the things that I want to mention that this one, also, is equipped with an ILM.

41:11

So the customer can be able to utilize that key features as well.

41:15

Um, power side I want to stop here for a brief moment, and then the goal very quickly on the power, actually is this particular tower that you are seeing is a not rack mount version.

41:29

The reason why this was available is because there's a customer schools looking for a non replicable system that can be able to put that into their office, is the order of the working environment.

41:42

At the same time, Bill wants to stop for a highest fertility CPUs, very little amount of the games.

41:50

But then at the same time, don't want to have a data bad would be a lot of a PCIE link coming out on the back.

41:57

So we are using this particular want to meet that requirement. And we have a customer who's looking through an office environment area in looking at this already.

42:09

Last one that I wanted to show is the Strategy part.

42:14

So, we've been challenged with the, a lot of a consumer grade people.

42:20

How come, we do not have a thread repo, which is the principal migrated, workstation type of rising CPUs.

42:28

So we come up with this server, which is the tower based on the ripper from the AMD side. At the same time, we can be able to do a lot of a PCIE link coming out of the back.

42:43

So this one, I can't really talk too much about but what's your C E is while we have inline to come up with a sometime in the Q one time.

42:52

That's all the slide that I have. Any questions?

42:57

Yeah, So, thanks can.

42:59

Just a reminder, everybody, if you do have some questions for Karen, go ahead and type them into the question box, as part of your console area. You can go ahead and type in your question, and we'll get that asked.

43:13

I think one of the things that I I'm always most impressed about is the number and breadth of server solutions that are available from Supermicro.

43:23

I mean, this is not a scenario where you just take motherboards and drop them into chase season, and things automatically work.

43:31

There's a lot of engineering and design and cooling, and all these kinds of considerations that go into creating these models, so, it's not just a matter of, well, we're just going to take this motherboard and put it into this particular chassis, to create a solution.

43:48

There's a lot of validation that Supermicro has to go through.

43:51

So maybe you can just for a second, if you don't mind, could you kinda talk about some of that validation that you guys do.

44:01

Yes. one thing I forgot to mention at the same time is that we do have a building block solutions as well, so we do have a ATX board and EA exports.

44:12

The way that we validate or got started people, we call into the validation.

44:17

The way we designed, the server is mainly the requests and the market trends.

44:24

And once we get those deformations and we put our heads together, and then they come up with an idea of what is the best solution you'd be able to provide to the end users.

44:34

And then, when we, when we come up with a server or when we come up with a design, we go through an engineering review at the end. We look at the Engineering View, along with the AMD ... games, the same thing.

44:49

The Intel based product lines and that, we have to get reviewed by the e-mail and the AMD side.

44:54

We have designers heavily dedicated for this particular product lines, that we have weekly meetings, and then, we finalized the design, what, can and cannot do.

45:07

And then, we go through the engineering view, printing up power, and we go through the iOS update, An IP, and I have a that we work along with the AMD to get the support and drivers that we need, epoch, the Windows side.

45:23

And then we go through the We power power ups always loads.

45:29

And then the EBT DVT, and then the ... and ..., and in the sample.

45:36

So all those times, the usually, a product lines, a kickoff to the final takes, it depends on the complexity of the server.

45:44

But it could be as easy as 3, 4 months.

45:47

But you can pick up, too, now, sometimes eight months, to a 9, 10 months.

45:53

So that is our ability, and because of that, GPUs are coming out very strong V space, we spend the majority of time validating that as well.

46:04

So these days, we anticipate to have a lot more GPU users.

46:11

So we, we take thermal very, very seriously, more serious than we ever had before.

46:19

So, some of the questions coming in. We have a collection related to kind of the design of the servers, regarding the A, oh, I am. Modules Have those.

46:30

Then designed so that you can replace them without having to remove the motherboard.

46:36

Yeah, so a IOM and SI O our modules are basically the PCI is aggregating feedback.

46:42

So you can't just the deepest. So let me go. I think I'm still showing the slide. Right?

46:49

You are, yes. Well, I see it, yes.

46:52

It's a good thing to show is the Bitcoin side. I think that Bitcoin has a bigger picture.

46:59

So it's a similar idea, as a S I, S I O N, So this is the.

47:07

Wait.

47:12

This particular area right here, that, I'm not sure if you guys can see my cursor, but the mic cards that you see.

47:20

I'll keep on losing. Sorry.

47:26

So the new cars that you see, the, or the ports that you see coming out of the bat.

47:32

Sorry about this.

47:34

OK, I guess in my cursor goes over.

47:36

That means it's kinda change the slides, so the word, if that's the key features, if you look at on that side, it has a four and nick ports looking thing.

47:46

So that particular area is a way of the SI O M, or the ... available. That's the way it goes.

47:56

Um, no, I'm sorry. I think a better picture will be.

48:04

My apologies.

48:06

Data picture will be this.

48:07

So that particular area it has a little thumb ***** where he can remove the IOM. And then you can just push the new It's like an PCIE AOC type of design. But instead of taking the top cover off and then put the PCI you to amount and whatnot, it's a very simple design where you just just unplug basically, from the mother-in-law and the Weave cloud backend.

48:39

And then the AI IOM, you see that ...

48:45

compliant, meaning that the off the shelf, at OSEP three point O version of the card can be populating into that box, as well.

48:56

So, now, my Maps already had I made an announcement that, if they do come up with the OCP type of cards, not just the PCIE.

49:06

So, they can be able to use that particular cards onto the hope that answers the question. Great, I think that's, that's perfect.

49:15

For some of the raid setup's software raid supported for U dot to NVME.

49:26

Personally, I have not tried it but we have tried the OS, raid onto the setup.

49:33

So let me pick that up on homework.

49:37

OK, So looking at the Redstone chassis and And by the way, we will provide supermicro with all of the questions that you guys are submitting.

49:47

So you know, Cal needs to get back with somebody on a specific question.

49:52

He can do that after session, through e-mail, or if you guys type in a question and we don't get a chance to ask supermicro, be able to follow up with you guys after the session.

50:03

So feel free to keep sending in the in the questions and we'll get them over to the Supermicro team.

50:10

But regarding the Redstone chassis, the two U Redstone chassis.

50:16

Are there any heat issues with that if you've fully loaded that chassis with CPUs and GPUs Now so all our validation I forgot to mention is that the we fully populate everything. Especially for the thermal. We use the highest GDP CPUs and then via we populate everything to a max.

50:39

So to prevent from the any hiccups that we might see from the field. So we use the highest and the greatest to accommodate all those possibilities.

50:52

No, no issues of bottled up to you to your phone factor on the restaurant.

50:57

All right, so before we get too far away from the the IOM question, I actually, I should have asked this because I don't know what does that stand for. Can you, can you tell me what?

51:11

AI O M is a advanced audio module?

51:17

OK, so, we don't, we want to be unit, right because we don't want to use the word or common knowledge that readily, you know, used by the, I should say, OCP community.

51:31

So E A or C P three, that will, kinds of things.

51:35

And, D, because the OSEP communities being pushing to have the specific cards validated.

51:44

And then that they've been battling with that particular community for, within themselves for quite awhile. And then the world does not really budget in oral or favorite their side.

51:57

So, I mean, as you all know, that OCS OSEP has been out in the market for, as long as I remember too, early, 2010, to my knowledge. But I know it's been designed way before that.

52:12

A lot of customers who are using the OSEP, it's not the same rack mount, right.

52:17

It's not the same use, use, dice and and that the height of the, you know, use a different than a regular rack mount. So, take off.

52:26

I owe you, instead about, you know, just the regular are all ..., so, 1.75 inches from the regular rack. You busses the OCP.

52:38

If I recall correctly, 1.889 or something like that and then the it's not really a rack mount for 19 inch. It's actually like a 21 inches or something something different.

52:49

So, a lot of, you know, OSEP communities, they want to, you know, explore the regular rat mouse ever changed, too.

53:01

You know, persuade them to move over to using an OSEP. But, but again, it's a battle between them and the nowadays that they want to have what is a specific car to be compliant and alumni, and, and, and now they finally come up with their own, you know, form factors of this particular design. So, that's where we stand as OSEP today.

53:23

Great, thanks. Um, so I'm going to try to squeeze in a couple more questions here and then, and then we're going to wrap things up.

53:29

But again, if we don't get a chance to ask your question directly, we'll make sure we get those over to Kim, Jong and team question. You guys have a lot of variety of servers. A lot of different server solutions.

53:42

not just AMD but also Intel are, do you have a configuration tool available that would help customers identify, you know, what, what products they should be using in an integrated together, any kind of tool or resource available?

54:00

That, actually, I don't know, I mean, because it's hard to answer where it might be better to go with an intel or AMD at the tool to decide, might. do you know if there is any tools that are available for them, or for the end users?

54:17

Not all of the sources that we have a product selection is within our reseller portal, which customers can register using?

54:27

my e-mail address and be able to select a product from there, but I don't think we have in terms of market segments. Like what what's the best fit for the ...? I don't think we have a tool.

54:45

We like to talk to the customer and help us help us select what are the best option that's available for our product. There's a lot of products that we have that are fitting into multiple market segments. And based on each customer's requirements or sales group will also appears FA's

weekend, we can talk to the customers and help them decide which one is best for their particular use case.

55:13

Great. So let me ask one more question, and then, and then Mike, I'm going to ask you to talk a little bit about the promotion you have, and then we'll get everything wrapped up for the end of the day and the week.

55:24

But, um, can on the high-end servers do those require 220 volt? or what's the requirement there for the high end servers?

55:35

So all our products are our power supplies are supported from the 1, 10 up to 240 volts.

55:45

So, it's in our standard to support the variety areas of the the customers.

55:56

So, it's a it's a power supply design rather than the server itself.

56:03

So, if you go to our, lets us a website, for example, landing on the specific server, either they add, So, let's go to either supported under the power supply section, OK, Great.

56:18

So, um, Mike, I know you have a promotion we want to talk about.

56:24

So maybe give everybody that quick update on that, and what's going on with the promotion you get.

56:30

Sure. Currently, we have a promotion running with ASI.

56:36

For a bunch of CPU plus server or CPU plus motherboard bundles, there is a rebate that's going on until the end of January four for both SQL socket sockets. So, for more details, please do contact your sales rep.

56:57

Again, it is a product selection, supermicro side. We love to talk to customers and customer select product lines. So, anytime there's, the question is in terms of what best fits the use case that surely server to use? Please do reach out to your sales rep and they can reach out directly to us. All the sales rep has direct access to civil micro sales reps here. We have a whole to support a size. We have a wholesale 5 plus 1 FTE.

57:26

And additionally, we need to, we can pull the idea of the pieces that, that, that are specializing in their products like.

57:37

Great. Thanks, Mike. Yeah, I just want to, you know, re-emphasize that, that, you know, the partnership that ASIS add with Supermicro is several.

57:46

As, you know, decades long. So, we've been a partner for a long time.

57:50

And the engineering and the resources that Supermicro puts behind, they're server products and making sure that customers have access to the right product and the right solution is really quite phenomenal.

58:04

They don't have just a few servers that fit a few markets.

58:08

They have servers literally for every kind of market.

58:12

You saw a lot of them today during the presentation whether it was a GPU server, dual processor server one, U two, you know, large, high-end servers.

58:23

I mean, the amount of product line that Supermicro has available, because every vertical within the server market is quite impressive.

58:32

And the engineering and, and design that they put into making these is really great.

58:37

And the other really strong advantage as Supermicro brings to all of us in the channel is the fact that they have these solutions available at launch.

58:47

So when AMD launches a new processor, when Intel launches a new processor, Supermicro has the platform and the solution readily available so that you can get to market quickly and leveraging that is really a major selling advantage. So, I just want to mention that and reinforce that here at the end.

59:08

So, it's one of the things that we've enjoyed about working with Supermicro over all of these years.

59:14

So, with that said, let me also kind of wrap this up. In general, I want to know, thank you, guys, thanks, Supermicro, for doing this presentation for us. We really do appreciate it.

59:27

So, Keon, Mike, thank you so much for all the customers that have joined us today, and throughout this week of events, we really do appreciate you participating and joining us for all the different sessions that we had, whether it was P N Y. Covering Quadro, or Intel talking about ... Pro.

59:46

And yesterday, Samsung, going over all their SSP Product Lines and today, Supermicro, educating us about the AMD Product Line. We really appreciate having your participation.

1:00:00

So there's a few things I just want to share with you guys do, look for follow up. E-mails from me, I'll be announcing a lot of different things related to prizes that we've given away links to the webinars that we've recorded, and other types of information just related to the sessions that we've had.

1:00:18

So look for e-mails from me and try not to bombard you too much, but, you know, if anybody is familiar with me, they'll tell you that I tend to over overwrites in my e-mail.

1:00:31

So, I apologize in advance for being kind of long winded, but that's the way that I like that, to type the e-mail. So, just look for some stuff for me.

1:00:42

We will be announcing by e-mail the winner of the Beats headphones.

1:00:47

And the 10 50 dollars American Express Gift Cards, as well as the Intel server system.

1:00:55

So, you will see e-mails from, from me announcing the winners of all of those products.

1:01:02

Just so that everybody knows, hey, who, who won these, and where they really, truly are facts given away. So we'll make sure that you guys all receive the announcements for that, and hopefully, within there, there'll be, no, some 12 lucky winners that will be celebrating some prizes that they can, they can get.

1:01:18

So, with that said, can Mike, is there anything you wanted to add at the end before we close it out and let everybody head off for the day?

1:01:29

Tom. Good. Thanks for the opportunity to share my slides.

1:01:34

Well, thank you so much, Mike.

1:01:37

Go ahead, Mike.

1:01:39

I hope everybody will have more of what we offered in terms of A and B servers, Any questions, please feel free to reach out.

1:01:52

All right. Well, with that, I am going to close that close out the session, and thank everybody for joining us on the call. You guys, Everyone, have a great holiday. Everybody, stay safe. And hopefully, down the road, we'll all get a chance to do stuff live again.

1:02:09

But we appreciate you joining us for this session and this event.

1:02:15

So have a great rest of the day, everyone will talk to you soon.

1:02:19

Thank you.