ASI Technology Summit Q4 2021
Supermicro X12 Rackmounts: High Performance Systems based on Intel 3rd Gen Xeon Scalable DP Solutions

0:05
Good afternoon everybody, this is Kent Tibbils with ASI, I want to welcome you all to day four of the Q4 ASI technology summit.

0:15
Today is our final day for this particular summit, so I wanted to thank you all for joining us for this entire week and for being part of this event.

0:25
We really appreciate everyone's involvement and participation.

0:29
You know, they've been very successful sessions and well attended, and the feedback has been really positive.

0:36
So we appreciate you guys joining us today.

0:40
Being a part of our event today for our final day, we have supermicro for us today

0:49
Yeah.

0:50
Here I hope I got Truong name, right. But Truong is here today, and he's going to be talking about the supermicro X 12 rack mount products.

0:59
But before I turn everything over to Truong, just a couple of things I wanted to announce, we do have some prizes from yesterday that I wanted to announce those winners.

1:09
And like we've been doing through all of the sessions, we have the big prize it Toshiba gave away yesterday, which is the 50 inch LED TV. And we're also giving away 5 50 dollars gift card.

1:24
So I want to give the gift cards away first we have for the $50 gift cards.

1:31
We have Speed Bowmen, Jason Lander's, Marks Dewey, Larry, Dan and Charles Zeller.

1:39
So congratulations to all of you, I'll reach out to you by e-mail sometime next week to kind of collect information for getting you your cards and for the Toshiba, 50 inch LED TV, we have promised fleece. So congratulations, Thomas, also reach out to you by e-mail.

1:58
So for today's session from supermicro, just to remind everybody, we're be giving away two power beats wireless headsets.

2:08
So there'll be two winners for that each winter.

2:10
It gets one headset so we're not doing one winner gets two headset, two winners get one headset, so we'll be giving those away. Since today is our last session, we will announce that the following week through e-mail.

2:23
Um, last thing before, we kick everything off to Truong , time to get warmed up and ready to go. Here are the last announcement I want to make everybody is that for questions.

2:35
If you have questions, go ahead and type them in to the question box.

2:39
I have myself and Shelley here today to help monitor the Q&A session, so you'll either hear me or Shelley monitoring the questions.

2:49
Just to let you guys know, we got a couple moderators on today for our our final session. So without further ado, Truong, I'm going to kick it off and hand things over to you to get started.

3:01
Alright. Thank you. Thank you. Can think of Shelley. Think of the audience. I know it's been a long week for you guys, but hopefully batting cleanup here, no pressure is on for me to kind of deliver. But anyways, let me introduce myself, my name is ..., I am one of the Solution Manager here.

3:17
So I've managed through a particular products on the X 12, this cloud, DC, our Ultra, and our hyper products, all three products are going to be on your isolate processor, and these are all Gen four.

3:33
Excuse, alright, I wanna kinda get that out a way. It's kind of reduce and confusion. So we'll go ahead and start with cloudy C then we'll move on to old tried, and we'll move on to two hyper. So, one way to kind of distinguish between all three is good, better and best. So, let's start with the good.

3:53
So, cloud is see is new for micro 4 X 12. The previous product you may have known is called WI. Oh, we have basically four SKUs in the cloud EC two series.

4:07
We call it Cloud UC because this is really towards geared towards data centers, data centers, are looking at cost saving, but also looking for compute on the compute side. These are dual processor.

4:21
We support up to 270 in each of the cloudy, see whether you're going to use a when you or to you. What's new for cloud you see this year is AI on support. I'll kinda go a little bit or a little more depth on the IM and how that works for, for networking.

4:38
OK, for memory wise, we support up to 16 dims. The ... can be equally share with Intel D C P M M, or ping them as they call it this year. For storage options, this a couple of products here. The one you which is the far left hand side, top left, there's the 120 C T N 10 are he knows that their orange tabs so, these aren't SAB represents NVME.

5:09
The NVME are going to be an option for cloud. You see, just to save costs and some customers, you know, they, they prefer not to have the mean, they prefer to stay with the sat or sass drives.

5:24
So, as you look at the motherboard, you notice on the picture there's a clear distinction between the left and right? The left a little bit, kind of, cloud out, and the right? It's clear. The reason we do that is to represent an equally balanced architect of the cloud. You see, a lot of our customers are demanding that.

5:44
Hey, you have do Processor, we want each of the processor to be mapped to the networking.

5:51
We want the process to be mapped to the rear IO, so we kind of listen and took all that feedback, and we provided them really well balanced configuration. So, when will you want to look at it as you could take a knife, split it down the middle?

6:05
You'll get equal share of the platform itself, OK?

6:10
Also, four for cloud, you see, we're offering, do MME, M dot to Onboard! So that kinda saves you adding another ALC to support that and saves you the front drive. So lot of these M dot to onboard could be your Your boot drive.

6:26
So little bit cost saving and size wise for M dot to use the terabyte.

6:32
Usually enough for the boot drive.

6:39
Let's start with, the first product, is the 120 CT and 10 are. So, what you see in front of you is the so, the motherboards all across, are the same across all the cloud you see. You will have here is 10 dries. Either be savvy sass MME, like I mentioned before. Enemy is going to be an option. We have cable kits and the cable kits will provide the back plane to be supporting the enemy. So if there's something that customers looking to at a later time, it's best to add it now doing the production because you do, we do not sell the cable kit separately where the customer can install that themselves.

7:21
Also, if you notice in the second picture down, you'll see some blue and red. So the blue represents CPU two and the red represent CPU one.

7:31

So I talked about earlier about regarding the bounce configuration. So here you'll see each of the IOM and the rear PCIE by 16 slots equally mapped to CPU one CPU too.

7:47
This is R 120 C dash T R, so this is a non enemy support. You'll notice that there is, on the top right, I don't know if you can see my cursor, but here, we still offer a optical drive, if a customer is looking to use an optical drive, so we have a kit for that. Risk Slots are exactly the same as the previous model.

8:17
This is our 6, 10 C T R. So basically for drives, a little more or less on storage. So anytime your customers looking for a little bit less storage or more compute, This is a great model to choose from.

8:34
Now we move on to the cloud you see, but the two you node to you as you get 12 dry base here. This is each of the dry baits are 3.5 inch drives.

8:44
We have these particular dry trace that will support these 2.5 and 3.5. So using the same dried tray it will support both sides. So that's kind of a nice cost saving there. You'll also notice the iOS, a lot more reale compared to the one you, but you also see that everything is mapped out pretty well with the CPU one CPU to.

9:11
Let's take a little little bit closer, look on the Iowa expansion. That's always big, I mean, the fronts are pretty given, but the rear. So you have to IOM these AIM, or, if when you're networking, you go from a one G, all the way up to 200 G. Your Slot 1 and 2 are by 16 foot high, half length slots. So, you can put additional rate controller, or you could put a 24.

9:40
Here are some of the popular option. we have configure out with a lot of our customers here. On the left-hand side, you'll see with the rate controller, one, with the fiber channel nick, and they'll using it to 4.25 G.

9:53
on my right hand side, a little more GPU intensive. So we could put in to Tesla to force for the cloud, the C So it gives you a lot of flexibility on how you want to configure the Cloud EC.

10:10
On the two you expansion, you'll get six instead of the two. The six are kinda unique because, we use, a particular table that allows us to create A by 16 slot and how we created by 16 is all about moving us, allow us in production. Moving the cable.

10:29
So you'll see in the in the configuration where 2 and three can be configured either by A or by 16. So, if you have a GPU card, like the B 8200, of these can be put into the aisles.

10:43
So, here is a better picture of the route configuration where I'm having here. I've got a V one, I get two V 100th on top of the raid array card.

10:53

Or, you could go with four Tesla T four.

10:57
So, I'll take a quick pause here to see the journey kinda questions you may have, or for Cloud EC.

11:05
..., a question that I saw on one of your earlier slides, you had something called Fill Silicon Boot of Trust.

11:14
What what is that? There's that on all of supermicro products or?

11:20
Yeah, that's a great question. So, the the security is It's really enhance now on the X 12. So, we what we use is we updated our BMC before we were using the ASA 2500 chipset for security. Now we move to the aes 226 on it, so the 2600 gives you a secure boot, also gives you a golden image. So the golden image is for some for, let's say, for a reason, the customer corrupts the the original file. So there's a backup file of the golden image and that could be re-installed onto the server, kinda gives you that security. This, there's a bunch of things involved in that security but all the security for the Extra Husband enhance with 2600 chipset.

12:08
Great, So, just to confirm that that feature is available on all X 12, server models', correct?

12:15
That is correct.

12:18
So, um, maybe, real quick, what's the difference between, you know, X 12, series motherboards, and X 11 Series Boards?

12:30
So, really big difference, because we're moving from X 11 motherboards are, or the previous Cascade Lake, user task, a lake refresh, or the casket like a regular classroom, like you're looking at little bit slower, in didn't size 2600, well, now the new extra support. I went to, 3200, you're going from X 11. You've gone from the previous gen three.

12:58
When, now, when you move to X 12, which is gen four, which is twice the speed, so quite a bit of a difference in, in technology, Performance from X one to X 12.

13:14
OK, So we're good to keep going, If you guys have any questions or additional questions, please go ahead and type them into the question box, though, Trump We're, we're good to move on.

13:30
Here's a little bit of a naming convention for the X 12, just just in case you're looking, searching our website, how you want to find X 12, cloud?

13:38
You see, this kind of breaks down in the, name it, in the name itself, and what does all the numbers and letters kinda stands for?

13:49
This is R 3.5 inch drive, so very, very similar. We have agus, pure breakdown.

13:55
There are meetings behind these, these model marlo names.

14:06
OK, so that was cloudy. See that we're going to move to Ultra. I mean, I know we push a lot of Ultra ultra, probably one of our flagship product here as well macro. So you want to think about ultra think is a bigger brother of cloudy seat. Remember we talked about good, better. So this is better. How's it better? It gives you a larger footprint. We were looking at 16 ... on the cloud. You see now we're pushing up to 32 dim on on Ultra. Alright? Processor supports can be the same all up to 270. 1 thing we did with this year is you'd notice all the orange tabs.

14:43
So, all of our Ultra has a hybrid battling, hyper battling, meaning that it also, it will support sada SaaS. It let me know. If you're looking for SaaS use, you still need to add HBA card or re controller.

14:58
For any me, very similar to cloud, you see where is an option, so we have cable kits and the drive trays available for you. So we have different, we won't kind of go into that. It's on our website, but as well, I kinda point that out. That enemy is an option, OK.

15:18
Also, what we did with X 12, this year, was we reduce the, I think, a lot of the headaches that X Levin was known for X 11, I think, has something like 53 different skews, which is astronomical in trying to manage and trying to educate the, all the cells for the internal and external.

15:40
So, what we did, by using, we've done away with a lot of the skews by doing a few things, one.

15:48
We do the hybrid back plane, so, before, you can say, Hey, trial, which one support for enemy, which ones have put six enemy. And, I had to give you a particular model. So, this one, now, you can say, Hey, I just need for having me, OK, I could go with a 120 you or we say, Hey, I want a 12 ME. I could go with the 620 you. So that saves us a lot of headache also for this year is the we change out the Ultra ... so ... gives you a few different things. ... gives you the internal, gives you the internal connection, but also gives you network connection. So one thing we did this year was we provided a ultra rise at the has no network. The kind of her love of the customer and say, Well, I don't want the the 1, 10 or 25 G on board, I'm already using a 100 G. I don't want to pay for that.

16:40
So for the Ultra, you could configure it with No, Nick, and you'll say, but you save yourself a little bit of cost. We'll kind of go through a little more as I'll walk you through all the products, but, and we'll see what the different stream from X one to X 12.

16:59
So if you, this is kind of a quick cheat sheet, is some customers asking, hey, try and what's between Islam and X 12. We talked about reducing that emerged skews. We got to carry the same optimize architect from the cloud. Is C over 2 X 12 ultra.

17:16
So more of a balance configuration, especially if you are adding a 100 G network card or adding a GPU, we want to maximize the dual processor on board. A CPU support, 270. We've talked about lose about other members support for moving from the different generation 2933 frequency was on X 11. Now you have 3200 with 32 dim slot. This will also support the P Now if your customers looking to populate with the PNM, I know ping them could be pretty tricky. But feel free to reach out to me up to help you with the pinup configuration.

17:55
There's a certain ratio that we need that kind of follow not to break the memory configuration from Intel for the enemy. So we have when you told that will support a maximum of 12 ME. And for the two, you will support a maximum of 22 NM meet. These are direct attached, that meaning we're not going to use a PCI switch in the bat playing any longer. Each of the drives have a dedicated lanes. so we dedicated for peace, a link for each of the Indian meet these all gen four, so you'll get the fastest response.

18:29
You get the lowest latency for this service security feature. We kinda highlight that regarding the root of trust and the sign I PMI, the bios. So this also support the aes 2800 so that your customers looking for a more secure server, as 12 as a way to go.

18:51
Little bit of a naming convention So what we did, if you remember in the previous X 11. We had something like a 1029 U dash to our 40.

18:59
So what we did was we, we took away the networking in the name. So now it's going to be as simple as a 120 ... or a 6 10 ... or 220 .... So that kind of simplify things from a naming convention and also from a network convention.

19:20
This is just to represent 3.5 inch drives.

19:24
I'll quickly pause there to kind of kinda kinda glad you guys catch up on the ultras there any more? Is there any questions on Ultra before I kind of go through some of the SKUs?

19:35
I had a question. You were talking about the NVME direct attach maybe on a couple of slides back. If you're doing NVME, direct attached, does that affect hot swap capability of the drives? Are they still hotspot?

19:51
It's still hot swap.

19:54
So, then, the direct codec just gives you better better connection, better performance for your Gen four trick devices, correct?

20:07
A view in a previous X It's all about figuring out how to share the number F one, You know, you only have so many lanes to share.

20:16
So what we did here is, we dedicate four lanes to each of the NVME drives, so they get maximum throughput to each of the drives, especially in our X 11 to you.

20:31
With a 24 enemy in the previous generation, where we had to use a PCI switch on the back plane, The kind of helped share all of the.

20:41
Helps show all the 24 MME drive. So this year, we've kind of done away with that. We have more beside lanes from the isolate. So now we're dedicating each of the drive have their own lane, So that gives you low latency and the best performance.

20:58
OK, fantastic.

21:00
And we're good to move on, OK? So, this is R oh one. when you 12. So, if you notice what one thing with this year is, we add two more dry. So, not only are we increasing density for capacity wise before the worst 12 drop in 10 dri based. Now, we have 12 dri based. All 12 can be sad. Our SaaS and then me, you'll also notice on the second picture down is the configuration on the rear IO. So, with the Ultra, when you you'll have for iOS there are 2 slash 3 and 4 are by sixteens. Slot two is A, by 16, but it's a low profile by 16, the Slot one, is your Internal slot, internal slots are? Normally we kinda the array controller there to give you a Rate, or HBA functionality.

21:51
I talked about live about earlier regarding the network options. So this year we still kinda carry over the 10 G 25 G, Tangy, SFP Plus, and do 10 G but we also offer the no net solution.

22:03
So like I mentioned earlier, if your customer has, I'd say what energy they want to use, you can say hey you know what, let's go with a known Nick option. So I'll save you a little bit of cost there.

22:20
Here is our layout and how we are mapping the drives to the CPUs. You notice a slot 0 to 5 are coming home CPU one. Slot 6 to 11, I was coming from CPU to. So, not only are we balancing the rear iOS, but we also balancing the front drives, especially when you're loading this up with NVME.

22:42
So a lot of customers, like, hey, I want a balanced configuration, I just don't want a more workload on CPU one and CPU to has been idle.

22:50
So, we kind of heard loud the customers, and we kinda responded with a new architect for X 12.

22:59
Just to kind of, we've touched on, the, the 120 you grew IO, so slot 3 and 4 are your by 16.

23:06

Now, there's still will accommodate, whether it's looking for 100 G card, or you could put in a one, do with full length. Like, say, a 100 card start to is by 16 low profile.

23:19
So, you can move that 100 G card into slot two, and they'll still be able to support not only a GPU, but also a really nice nic card. Slot one is your internal that is for raid on HBA car. Or you can even put in a ... like that, like ..., To ..., Option zero. Is your authorizer authorizer is going to be carrying our network connection? Whether it's going to be the 10 or 25 G or even the, No, no, not at all.

23:55
Will just quickly jump onto the 220 UT NR.

23:59
You want to think about Ultra from 1 U 2, U two distinct as supersize files, right? You got the, the one you. Now, you're going to move to you say motherboard.

24:08
Well, before the one you had for rails, now you have eight were aisles. So a lot more things you could kinda cramming here looking for where you're looking for more drives, more IO support looking for.

24:23
Say for, for example, art or chassis here, mechanically we could fit up to four double why GPU that will kind of take that with a case by case just because thermal is always becomes a hot issue here. So if you have a configuration that has four GPU, whether it's going to be a one hundred's or the forties. And you'll pick them a high processor. Please let me know, I could work with you directly to kind of figure out if we could support that particular processor with the GPU. But mechanically with good fit for for networking. We're going to carry the same over two from the one you to the to you, whether you're looking for 10, 20, 5, or 10.

25:04
We also offer the no next solution. So, lot of flexibility in the two you compare to the when you. But the two, you will give you a lot more, a lot more flexibility.

25:18
Also, by the way, the motherboards are exactly the same. So anytime a customer's having managing their own data center, but that's one picture could tell them, Hey.

25:27
Whether you're looking for ultra when you to you say motherboard. So, less inventory, less training.

25:37
Here's a kind of a layout of how we map the front drives to do the CPUs.

25:48
I think the biggest thing is the iOS for Ultra.

25:52
The design of the ultras, it gives you a modular architecture design. Now you'll see that rise a card on the right hand side. Let's kinda highlight out.

26:01
Now, by default, Slav, 5, 6, 7, 8 or by eight?

26:07
Like you could say, hey, I want to, I want to buy 16 there to support a, whether it's going to be a nic card or a GPU.

26:15
So we have a kit that will replace the Rise A card, that will convert 5 and 6 and 2 by 16, 7 and 8 and 2 by 16.

26:24
Same server. All we do is check out the riser and you'll get the support of that. So I think I want to kind of highlight that.

26:30
I think that caused a lot of confusion on someone's look at the website. Go, Well, I only shows by A How do you get about 60? So this is how we get it by 16, is replacing the internal riser, and this is all laid out on our landing page that gives you the model number and everything.

26:46
So here I kinda just kinda did the same thing but wanted to just highlight between the dip between the left and the right. So the left is the changing out the riser and the right is the default.

27:03
I'll take a pause there, kinda maybe answering the question that you might have on unchanging this internal riser.

27:11
So we have a couple of questions. Not necessarily specifically on changing this specific riser.

27:19
but when you're moving from, you know, model within the different models and you're adding, know, all the different kind of devices, you've loaded everything up with hard drives and your maximum configuration.

27:31
Is there any concern about heat inside of these systems?

27:37
Know, we are thermally everything's been validated at full, full, full population. So, that means we put in 32 dense, and we fully configure Ultra when the what's going to be when you are to you. So, we maximize the population, then, we put it into a thermal Chamber.

27:54
So everything past flying colors, the only caveat is when you start using GPUs, these are 300 watt GPUs, are run pretty warm.

28:05
So that is when we get a little bit tricky, where, hey, we asked you to kinda reduce the the wattage on the CPU, whether it's going to be down to 250 saved, the 270, or even down to 205.

28:18

When you, when you putting, let's say, for example, you're putting 2 a 100, will have to kinda look at and go, Wow, let's see how many drugs you're gonna populate. We'll see how much memory isn't there. So, we kind of take that case by case, but, but normally on a non GPU, you could fully map this out, and won't be have any thermal issues at all.

28:37
That's I think. the important thing to also mentioned and remind everybody is that supermicro designs and develops and engineers, these products, from the ground, from the ground up.

28:48
So there, doing the complete design and engineering of these chassis and motherboards together along with the cooling solutions.

28:57
So they're engineering something that they've created all the way from completely from scratch to make sure that not only does everything fit properly, but the air cool, and everything cools properly.

29:09
This is Compared to doing, you know, maybe a server solution where your, you know, building your own product and utilizing you know somebody's motherboard and somebody else's chassis.

29:22
Those combination of products haven't necessarily been tested at the same validation level as something that's super micro is doing, where they've designed and engineered it completely from scratch.

29:34
So, I think that's definitely important differentiation to make when we're talking about cooling and design and those types of things.

29:43
So, do other questions here really quick and then we can kind of move on But are there any options for fiber interfaces with the networking?

29:58
Are you talking about a particular networking card that needs to be supported? Is that the question?

30:03
I think the question is is that you are showing a lot of the different you know, adapter Board and the different connections are there any options for fiber?

30:16
Let me double-check.

30:18
If you could just send me an e-mail, I'll definitely kind of follow up with some other PM's on that one.

30:24
OK, so Yeah, we do, by the way, just to mention to everybody on the call, when you do submit a question, we do have a copy of those questions and are able to provide them to the presenters, So

we're not able to get to the question today, or, It's a little bit specific or technical in detail, and we can't answer it.

30:43
Today.

30:43
We are able to provide that to the manufacturers so that they can follow up and and get back to later.

30:51
So, couple other questions here.

30:54
Really quick, Truong, on the rizor card that we're looking at right here on this particular slide, Does this come with the systems, or this is something that you have to purchase a separate?

31:08
This is an option, so you can purchase this.

31:11
If you need a, Slot 5 and 7 to B by 16, we keep it as an option, just because some customers don't only only only going to occupy those slots with by eight, and there's no need to include a higher cost for a ride that they may not use.

31:32
OK, um, one other quick question. And then we'll, we can go ahead and move on.

31:37
On the, uh, ugh, really, the Act, well, S, TH that chef was that available yet?

31:51
Have a double check.

31:53
Is that, I don't know, I can't remember which server line that's from, S S F.

32:00
Is that, that's not definitely not an ultra, or cloudy, see?

32:04
The dash H, or F F is X 12 S, TH dash app.

32:12
Yeah, that might not that, that's a product line.

32:15
I do not, not support, so, but I'll definitely bring in the right PM for that one, OK, so yeah, great example of another, you know, question that we can definitely get to Supermicro, and we will get the answer back on that.

32:31
And I apologize for not asking the question correctly. So sometimes with questions, not being able to answer, actually my fault, because I'm not capable of asking it properly, but let's go ahead, and, and we can, we can move on.

32:49
No worries.

32:51
All right, see here.

32:55
So, kinda keep, kinda keep kind of keep going on the, the, the rails. I think it's very important for Ultra me, I think not too many servers and the two form factor will kinda give you this many options. Option could always leads to headaches but that's why we have PMS here that kind of help you navigate the ultra maintenance. You could throw a lot of this ultra and they will support a lot of a lot of things.

33:20
So I just want to kind of highlight the eight aisles of what they can do to, to support your customer requirements.

33:33
Now, let's move on to the 2 20. With the enemy, this is where it gets a little bit tricky, so I'll try to explain it the best I can.

33:41
For X 12, we do not offer 24 enemy solution, we offer a 22 ND me solution.

33:50
The 22 MME Solution does not use a piece, a switch or a switch in the back pain. And so each of the 22 enemy are Direct connect, so that gives you low latency and best performance. You also notice in drysdale one, that's kind of grayed out, like I read it, but it's a Burgundy tab. So that is your sat and sass.

34:13
So 0 1, a lot of our customers will put our boot drive, their connected, the onboard gives you a raid one with using the motherboard on Boot. Or you could put it to array controller and gives you a reward for your OS and from Drive 2 to 23 or you're an Emmy for your data data drives.

34:35
Now here's where it gets tricky.

34:38
Because of the 22 MME, all Direct Attach Whee, have our retirement cards.

34:45
You see Slot 1 2 for you See, the ALC SL G for the ALC software for E 40.

34:55
These are retirement cards, basically, retirement cards are pulling the pizza lanes from the motherboard to to the the backlink. So, that's how we achieve the direct attach these. If you are ordering a 22 FA me you have to be mindful that slot 2, 1, 2 and 4 are being occupied by the retirement card. So, you have to kind of look at from slot three to slot 3 or 5 to 8 are still going to be available.

35:30
So this is a no Nick solution.

35:33
So the know Nick solution, you'll get a slot three, still be available for, let's say, array controller, but if you go with a onboard knick for the authorizer, that slot goes away. So slot three now occupies retirement cards.

35:51
So that gets a little bit tricky, but on our website, we have kinda spells it out on which cable kit you need. And if you select certain cable kits, it will automatically populate the right Retirement card for that, so you don't have to kind of figure that out. That's all done on our website.

36:08
Soon as you select 22 MME, it knows to go, OK, whether it's gonna go with three retirement cards or the four Retirement Cards.

36:17
I'll take a quick pause there to see if there's any question regarding the retirement cards.

36:24
Not specifically regarding those cards. So I think we're good to go.

36:34
All right.

36:36
The last product on the Ultra is, are too short depth.

36:39
So, we kind of built this particular product for the telco industry, telco industry, have certain different requirements. From the regular industry, those requirements, a little bit more stringent.

36:54
one is, that certification lab certification is sort of a set of certification that needs to be completed and pass an order for certain servers to be selling to Telco, where there's going to be a Verizon, T-Mobile. Those requirements are liberal stricker, where they will put this particular service through drop tests, earthquake tests.

37:17
They will actually burn the server itself, make sure that all the flames are contained inside the survey does not spread. one thing we also notice here in the front are a little bit less dries, but the fans are actually in front of the server now. So you have six hot swappable fans, little bit less on storage. You have six storage, dry trace that, could be either soft SaaS enemy.

37:43
You don't, don't take the short depth as a little bit less feature Is. Still carry the same motherboard still, carry the same rails just that we shorten it up by six inches by moving the fan and do some space saving. Most of the these Telco data centers are not your high-tech data centers like you see here in Silicon Valley. But most of these data centers are near the radio tower. So these radio towers are often remote, little, more space is very constraint, and they run a little bit more, a little bit harder. So that's why you see the 6000 front versus the for fans that normally you will have in the, the, the chassis itself.

38:26

So we are working on that certification, and we expect to have this to be. Net certified by, I believe, by the first quarter of next year. So if you are?

38:52
Um, China, I don't know if it's me, but I lost audio.

38:57
Yeah, large body of two eyes, it's waiting to see it, if it came back, but.

39:05
Turn, we kinda lost the audio.

39:10
Are exactly the same as the previous 220.

39:13
So it just, everything is exactly the same, it just would shorten the chassis by six and just to kind of save the spacing on that.

39:24
OK, last and not least, let's move on. I am pretty, pretty late here, so I'll try to speed this up. So this is the external hyper so we talk about good, better. This is the best.

39:37
So, think about Ultra with all those features.

39:41
Now, hyper gives you all those features, but then some.

39:46
So hybrid just came out this year. It's still based on the X 12 isolate processor.

39:53
So what we've done this year with hyper is a few, a few things that a lot of customers are asking for.

40:00
one, lot of customers are asking for front IO.

40:03
So, the one you, hyper here is going to give you an option for, reduced the number of drives, but we also can support an ILM in front. So now it gives you not only the written network, but also the FET Network. As you can see here, will still under development for the two you, 24.

40:23
That's still moving pieces. But right now we have the 112 and the 2 to 12 available.

40:31
So this is based on the same form factor. Made the same spec as X 12, Ultra.

40:38
So basically dual processor 270, 32 them but what separates the hyper from the Ultra is kinda go into a little more details here is, one, we've done away with the screws. So fully tulis chassis, everything's Autumn Screws.

40:57

Lot of a customer has, you know, kinda hit us Hey, you know what?

41:02
We want to serve as a sales but the screws are becoming a factor.

41:06
So with hyper a full, all told us design.

41:11
Everything is ... from the motherboard to the back plane.

41:14
To the ILM.

41:16
This still has the 2600 BMC as previously with the cloud you see at Ultra. Right.

41:26
So, this is the hyper when you and to you form factor, So the one you pretty straightforward, where you can still get a maximum of 12 dry base, or you can reduce it to eight and use the dry bays, the four driveway to the right as a front IO, whether you're gonna add a network card.

41:45
Um, for the rear, we change it up a little bit on for the hyper was support three double wide, full length GPU cards, and we're using the AIM for Network. So we've done away with the ultra ultra riser and we move on to a I L M for networking.

42:08
This is kind of a breakdown of the the Hebrew when you itself, but he's pretty self explanatory.

42:20
Kind of the inner workings of the hyper, would you see the three do Dems?

42:24
You'll see the, the piece, IE, Gen 4 by 16 for where they're looking for to add a GPU there.

42:37
So, the ..., when you, you'll see from the far left, you'll see the the Standard configuration will give you the 12 dry base to rebuy Sixteens and the AIM.

42:47
For the one you, in the middle, was still under development, but we're looking at adding a possible adding 8 dot two plus a E one dot S for storage. On my far right, we're looking at six a U dot to enemy in the front with two ... for networking in the front. So, like a lot of flexibility with hyper things are still moving, but hopefully this will get released very soon.

43:16
This is our hyper to You.

43:18
Looks very similar to the previous with the Ultra, but on a on the hypervisor tool, is designed, get the same flexibility.

43:26
But so this is kind of the future for, for small micro.

43:33
We also have a two short depth, how I mentioned earlier with gardening, ultra short, short, short, depth. Now, we have a hyper too short depth, so.

43:42
also, going under level three, sort of occasion, or four, the telco industry.

43:52
This is the rear I also a lot of different configuration. Whether you're looking to add GPU support or are you looking to add more rear drives to the to?

44:01
The rear gives you a lot of flexibility on not only expanding the front, but also expanding in the rear.

44:09
Let's quickly chat about the alum Network here.

44:13
So, this AI network is basically leveraging the OSEP ...

44:17
Design, so this is a sort of micro proprietary design, but we're using the standard design concept. It's all the screws. So, basically, this gives us better thermal compared to our previous rear ultra risers or rear networking car where it might block airflow. So, this provides better airflow, and you can get this from one G, all with 200 G networking.

44:45
So, here is the LM networking. A lot less for, so macro, this all been validated. If your customer preferred to use their own, along as OSEP throughput, three dot O compliant, it will fit into also with micro.

45:04
Quickly highlight the MME parts.

45:07
If your customers are interested in what is actually in our enemy kit's, this is a breakdown of what cables included. The part number itself and the tray that is needed to make the MME option possible.

45:22
This is the same thing. But for the two U, 2 U 20 ultra to you, kinda gives you the same layout, Which of the kits required? What slots takes up for the retirement cards? And all the drive location for the enemy.

45:36
So, basically, a lot of customers are asking, what's in these kits? You could tell them, Hey, these are the cable that's in the MMU kit.

45:45
I'll try to wrap this up: EMC storage. With X 11, move to X 12, we are switching over to the new Gen four HBA, and re card, the, 38 0 8, or the 38, 16, or you, HBA cards? The 30, 908, and the 39, seeing very nice 16, or your rate controller. So.

46:08
Last but not least, I know we've talked about three different products and they, they quite a bit information.

46:14
But, this is a great slide to kind of fall back on, gives you a breakdown of the product compared between the cloudy, see, ultra hyper it gives you chassis depth number dri based.

46:26
IO support basically, everything you need on one page.

46:33
Workload recommendation. I know a lot of our Ultra hyper cloud is you can play in multiple verticals but.

46:40
We kind of put this talk together. Kind of give you guys a bit of where, which, which product kind of shines more uncertain vertical markets of this kind of maybe help you stood in the right direction.

46:57
OK, that's it for my side. I'll go ahead and take any questions that we have some time left.

47:03
Great. Thanks, Truong, really appreciate it. Thank you for running through all of that for us. A couple of questions here.

47:12
First, regarding DDR five Memory, and I don't know if you're able to speak to this or not, because we're kind of looking out a little bit into the future. But, as you look at roadmaps for supermicro products, what are you seeing?

47:25
In terms of DDR memory, and when, when do we think we'll see servers that potato that support that, that memory choice?

47:34
So, DVR five is going to be, it's gonna fall on the ....

47:38
We are in development for X or 13, So that will, for the four cloudy C N and hyper, those, two pollock will be moving forward. For X 13, you'll carry the same out of the dems 16 and 32, but you'll get faster DDR Memory Support on X 13.

48:01
And that would be four, Safar.

48:06
OK, so we didn't touch too much on sort of support programs or other things that Supermicro has available in terms of like transition maps and stuff like that, but maybe you can talk a little bit about supermicro's warranty and RMA Process or support structure.

48:30

So I'm not the best person for that, but I'll go and try it. Try my best I lead to, the Army support structure is.

48:39
You still have to kind of open up, file, a ticket that So that ticket kinda goes through that. The process itself as a system PM.

48:47
We have the authority to sort of provides either it's going to be across shipment or approval for the REM, RMA. And we have several different support structure for services, whether you're looking at a four hour support, 24 hour support, or 72 hour support. And that could be all laid out.

49:04
with the, with the particular pricing model you're looking at, or, Or requirements from the data center for the customer, OK.

49:13
So looking at all of the products that you showed today, there's, there's definitely a lot a lot of transition. I mean, you talked about Cascade Lake. There's like, there's going to be sapphire rapid coming. So all of these products are going through different transitions.

49:28
Does supermicro have resources on their Website for, you know, this model is transitioning to this model and sort of guidance to help customers navigate that.

49:43
Right now, we're not publishing it on a public domain. But we do, have some slides available that I can, definitely share with your, you and your team. That kinda gives, you live the direction on how we're going to move forward from X 12 to 13 just because we are under the embargo from Intel.

50:03
But my lot of my slides, I, have, are, would take the entire part out, but we could control the system itself. The kind of shows you, what's going to be on the horizon for 4 X 13, and what can you expect from the X 13.

50:19
And then what about like current models, what about from X one to X 12 with that?

50:24
Are there are there things for that or should a customer reach out to their ASI salesperson to kind of help them?

50:31
You know, determine what's the best model moving from X 11 to say X 12?

50:37
Yeah, actually, I didn't go through that, but I do have some slides available where we sort of kind of break down between how, how we can get a customer from X lambda X 12. I think one of the biggest thing is, is selecting the right processor. So, we have some benchmarking numbers that we provided that we can provide that shows you if you are, let's say, a 646 Cascade Lake. You want to move to isolate. We have a Corvette isolate. That shows you this is the best model. If

they want to keep that same frequency number, of course, this would be the best model move forward with with Isaac.

51:18
Have a question about add on cards specifically, so I just want to make sure that I read this question correctly, but the customer is asking if there's an option, four P O E, works as an add-on, and if so, what is the maximum number?

51:36
And as I asked that question, Shelley, can you take over the Q&A from from here, for me?

51:44
Absolutely, OK, So go ahead.

51:50
So it sounds like the cut, the side of the questions, is regarding how many add on cards can we add, too, to our server, is that correct?

52:04
Yeah.

52:05
That's the way I'm reading it OK.

52:09
So will start with the will start with Cloud EC. So cloudy, see you can add to to the the one when you form factor, and you can always add up to six for the two form factor.

52:25
For Ultra Ultra, when you you can add up to three or iOS.

52:31
for the two, you do that add up to seven, for hyper, that's going to be the same. Same number for as Ultra.

52:42
Great, thank you.

52:44
I'm going to take a stab at this one. It might be, asks, me read this one off for you. Will there be a focus on Level six, barebones for the X 13 generation, or is there more of a move towards fully integrated solutions.

52:59
I sort of mcelwee, we will go towards a fully integrated solution, just because we can fully control and fully validate all of the, all the hardware configuration, let us go into the server.

53:16
If you have this particular case where there is a fairly good volume, we definitely could take that case by case offline to talk about whether a barebones may may be approved for management side to kinda fit that particular project.

53:37
Thank you.

53:40
General questions I've seen a few of these commanded to about product availability, is, as she knows, there's so much inventory constraint. I don't know if you want to take a stab at that or if I should just kind of bounces questions back to an ESI person. But there has been a few that popped in about just the status of products right now.

53:58
Yes, it's, I think, Ed was filling the worldwide pinch regarding global shortage, is. we're not immune from that. It's always best to come and reach out to sumac yourself because these shortages are changing by the hour.

54:14
Um, we have a lot of OEM customers that all kinda fighting for the same piece of the pie, so it's best to kind of figure out what the project sizes weather demands are.

54:29
That way, your sales rep, as Michael can basically set aside the right amount of components to fulfill your order.

54:40
Thanks. I want to add on there, too, that will get this information out to everyone, but it does have a really good Supermicro support team. So we do have a lot of resources committed to this line. So I'll get to those e-mail contacts encrypts, because again, we have an internal team here, to help you as a first line for the Supermicro product.

54:59
Yeah, that's perfect.

55:01
Yeah, I'll just say, that's kinda my last question. I've seen come in, We do have, I know there's other questions. Yeah. No, we haven't addressed, it seemed a little more technical, behind the scenes, can't spend flag and things, and we'll take those offline and get, get these answers for you.

55:13
Mmm hmm.

55:15
And with that, I get a laugh, because I get a maintenance issue, so he jumped off service. I guess just shut up at the helm.

55:25
So, I just wanted to kind of recap today with, everyone is, today, we're gonna give away on a couple of super Micro, putting these out here, but we've got to beat Ken Is going to do the wrathful after hours, and then he's gonna, like, every day. I think he's been sending out a recap.

55:41
With the presentation slides, and a link to the video. So he'll do that tomorrow, and you'll also announce prizes for the grand prize along with the Supermicro beats. I believe. That's it on our end. Chang, Did you have anything more you wanted to add?

55:58
No, I just want to thank you for, for hosting this.

56:02
Thank you for the odds, for giving me the time to pitch the three different products. I know it's quite a bit to digest and 45 minutes, but you have any questions. Feel free to reach out to your ASI rep and they could reach out to me for any technical discussion you'll want to have.

56:19
Now, again, thank you so much for doing such a good presentation for our customers. And thank you all for joining today, and every day, this week, we've, we really appreciate your support.

56:31
Again, we'll get you all the information out, so if you have more questions, reach out to us.

56:35
And thanks, everyone, with that said, I'm gonna go ahead and end it, and thanks again, everyone, like a Chalet, bye!

Transcripts are automatically generated