Flash Memory Summit

INVT-202-2

CXL's Memory Form Factors: A Powerful Set of Options!

Presented By: Christopher Cox

Vice President of Technology

Montage Technology

35 YEARS OF NAND FLASH MEMORY INNOVATION

Christopher Cox of Montage Technology Presents…

The Wild West of CXL DRAM Memory Form Factors

Agenda

- Quick look at some of the new FF's
- Deeper dive into Form Factors options
- Varied Routing Solutions & impacts
- Packaging options
- Call to Action

The Gold Rush of new Form Factors

The industry has needed a new frontier for years and CXL is set to allow this new reinvention and mass expansion of form factors and solutions.

DRAM citizens (solutions) have lived in a city (very standardized environment) for years, governed by a host (SOC) that always had their best interest at heart, but will CXL DRAM Devices be able to handle the lawless west?

Image courtesy of Pixabay @ Pexels.com

Before…

What used to be a standard DIMM (Dual Inline Memory Module) with a couple flavors depending on if you were dealing with a Laptop, Client Desktop or Server system…

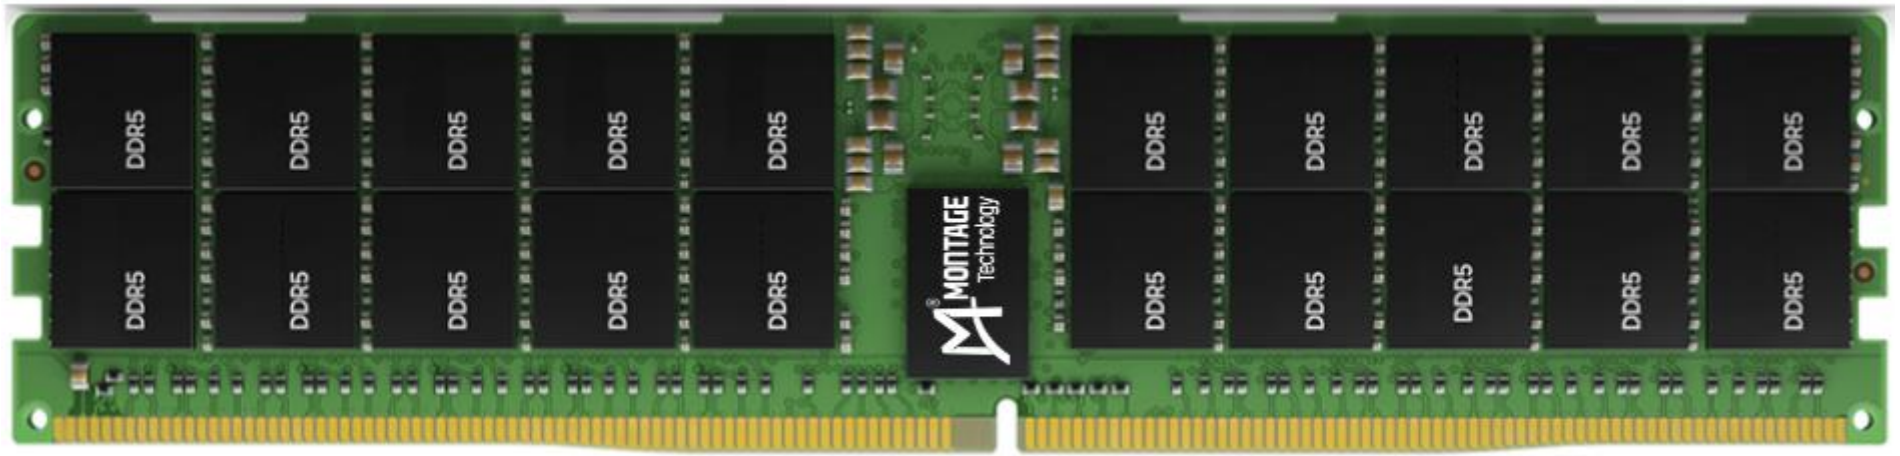

Has now exploded into all kinds of possibilities

EDSFF E1.L Formats

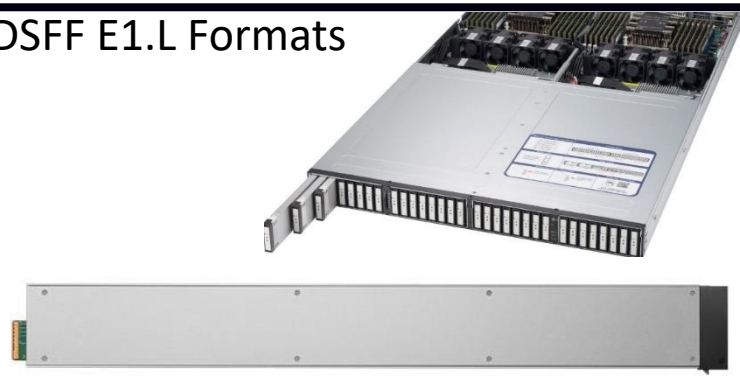

Custom Backplanes Chips + DIMM Slots

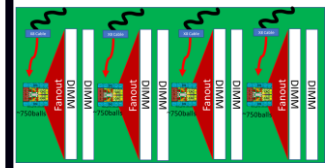

M.2 Formats

EDSFF E1.S Formats

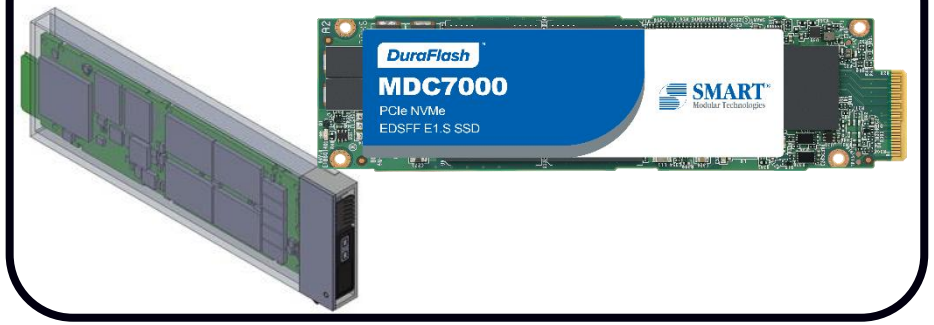

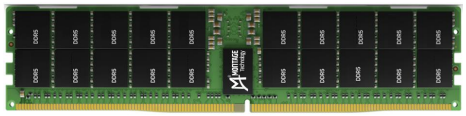

U.2 Formats

EDSFF E3 Formats

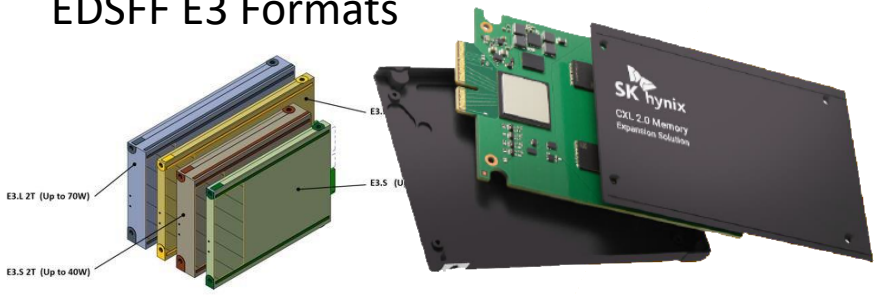

Nested modules in modules? i.e. GenZ FF to M.2

Add In Card Formats

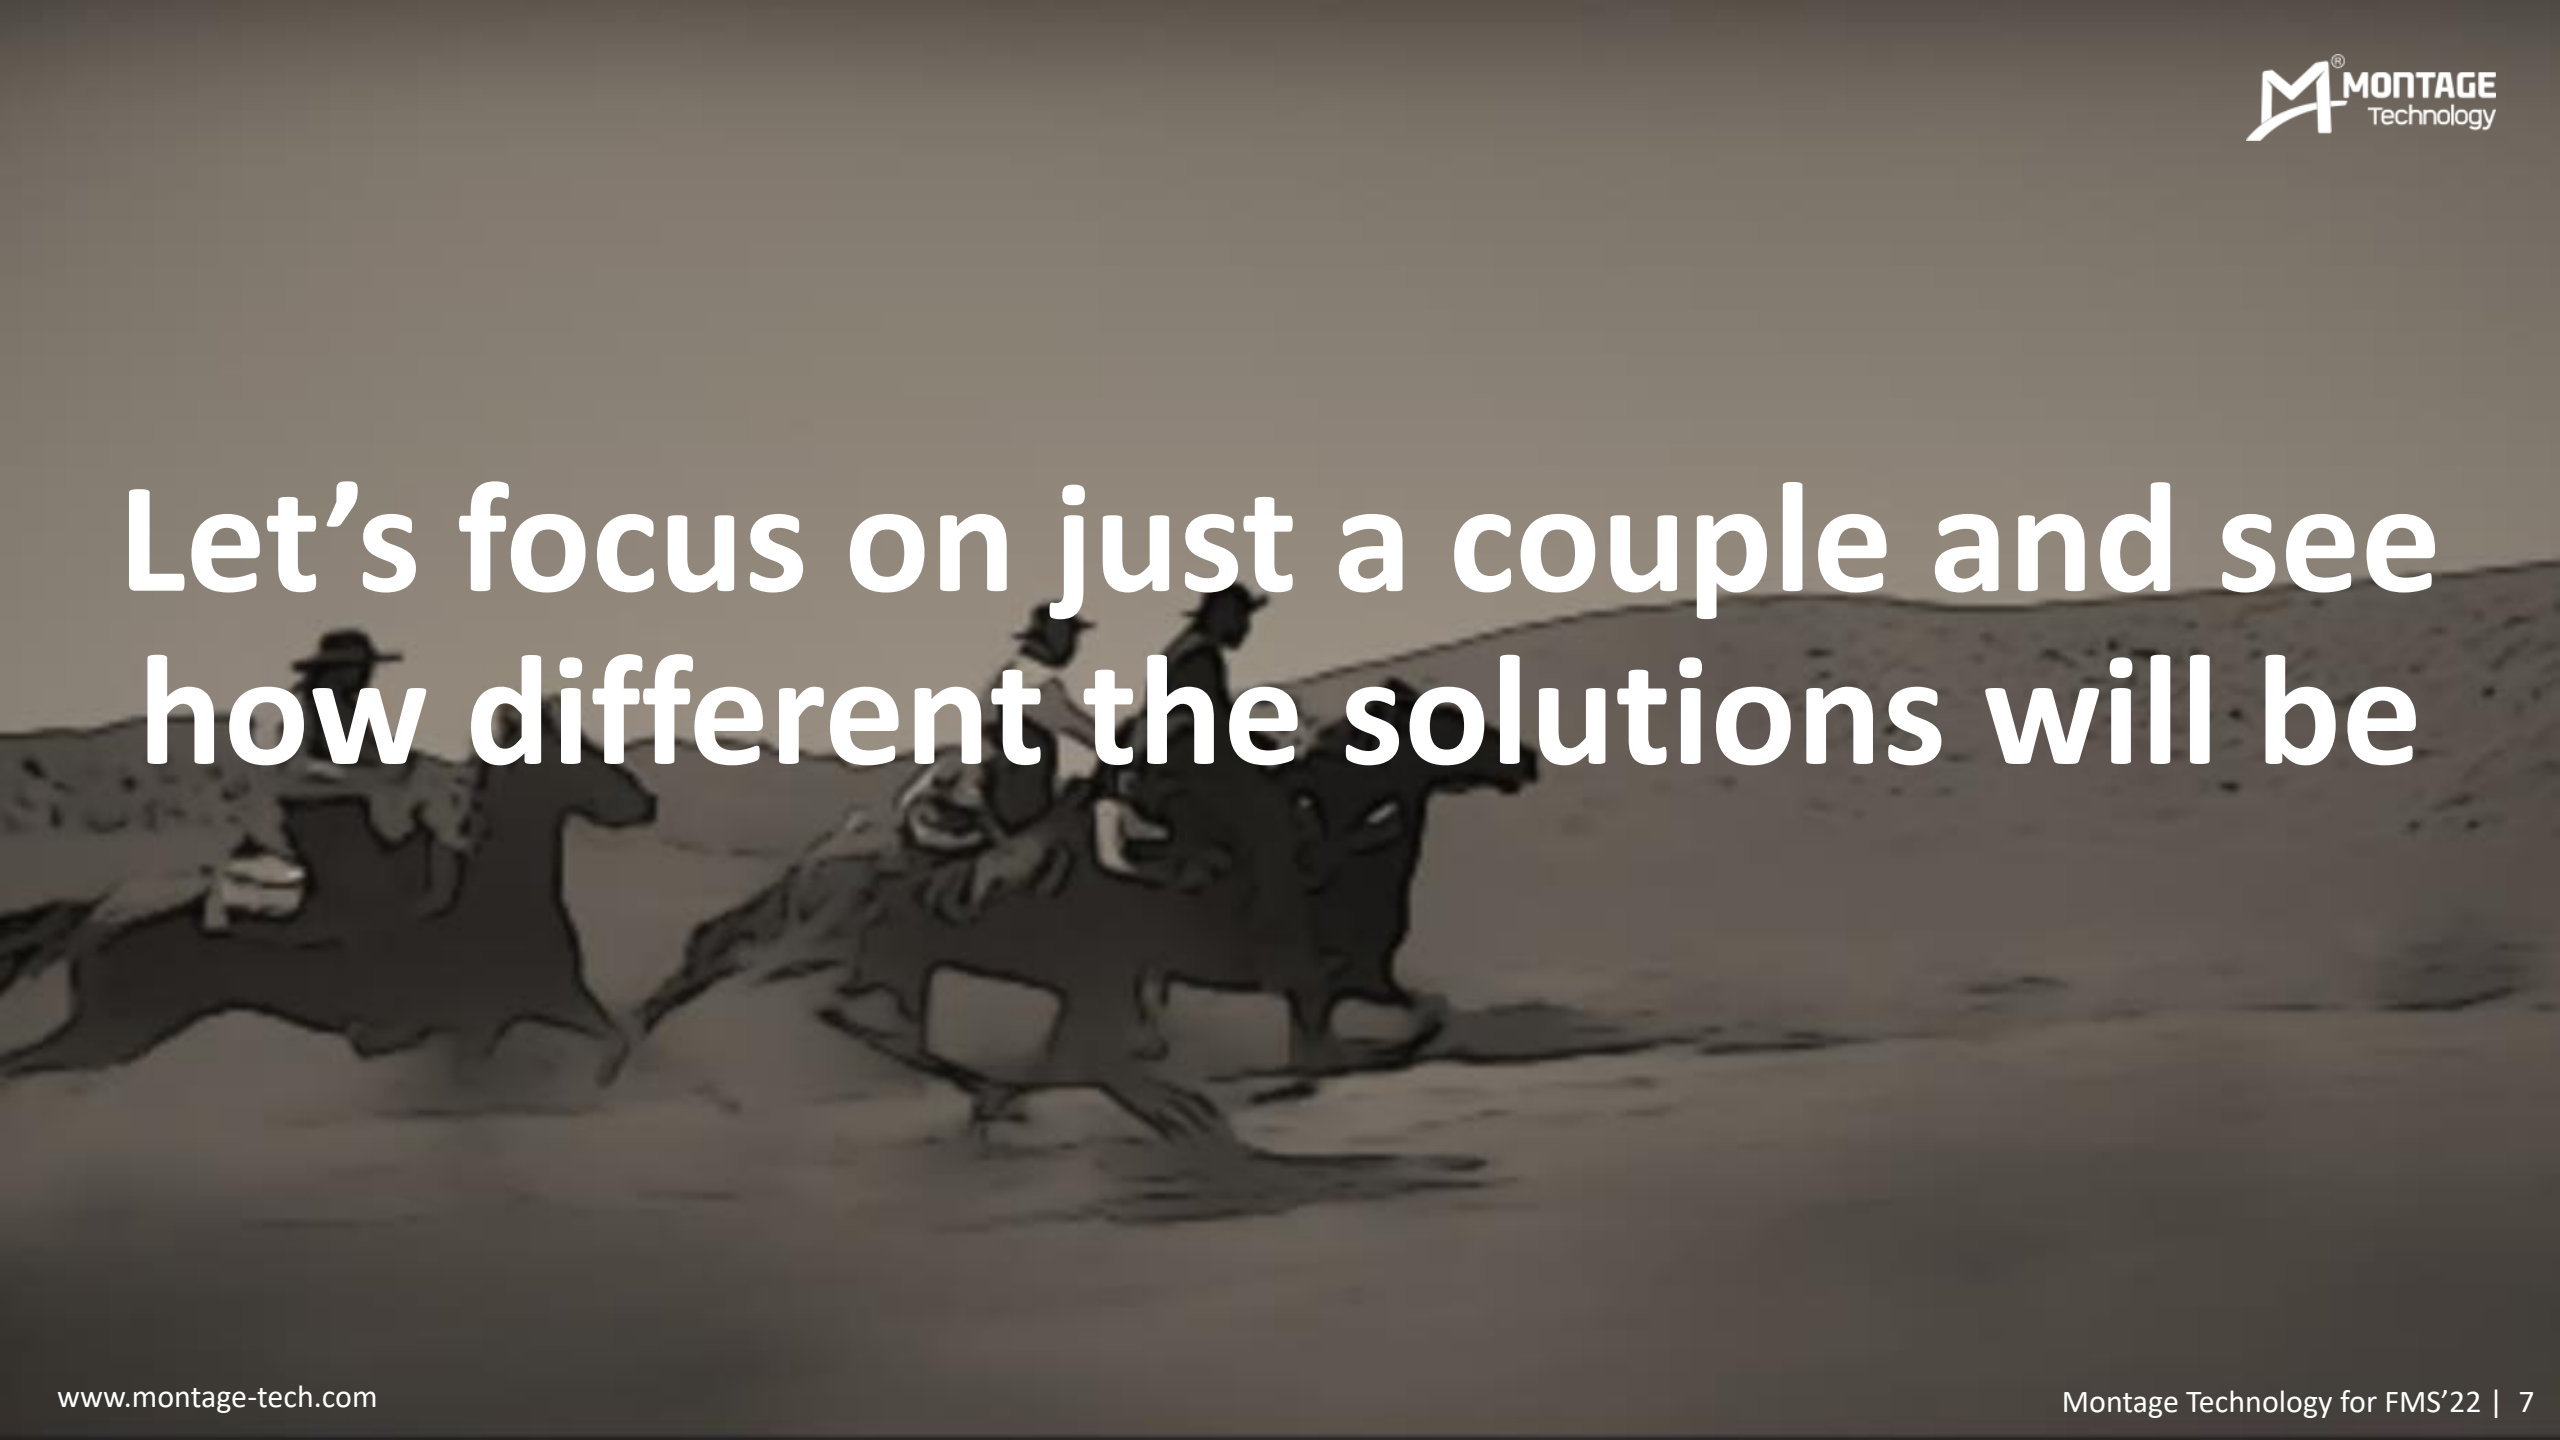
Let’s focus on just a couple and see how different the solutions will be

AIC w/ DIMMS vs E3.S vs E1.S vs DIMM

RDIMM

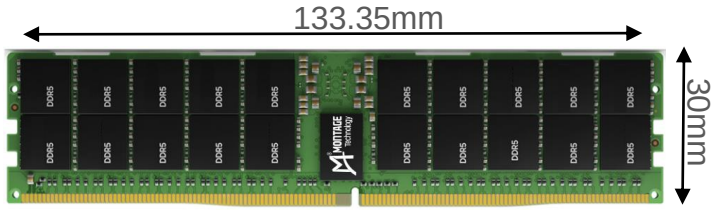

Thickness~3.5mm

SNIA EDSFF E3.S

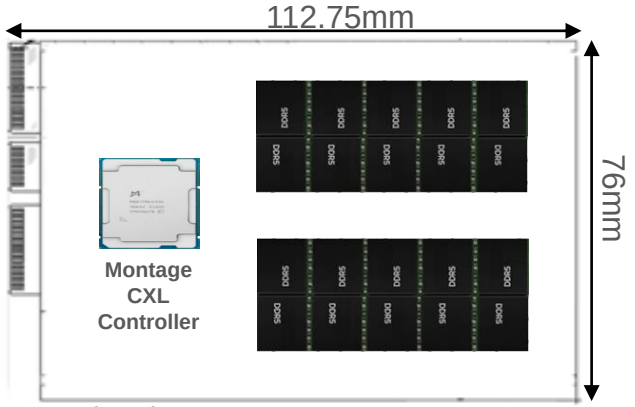

Thickness 7.5 to 16.8mm

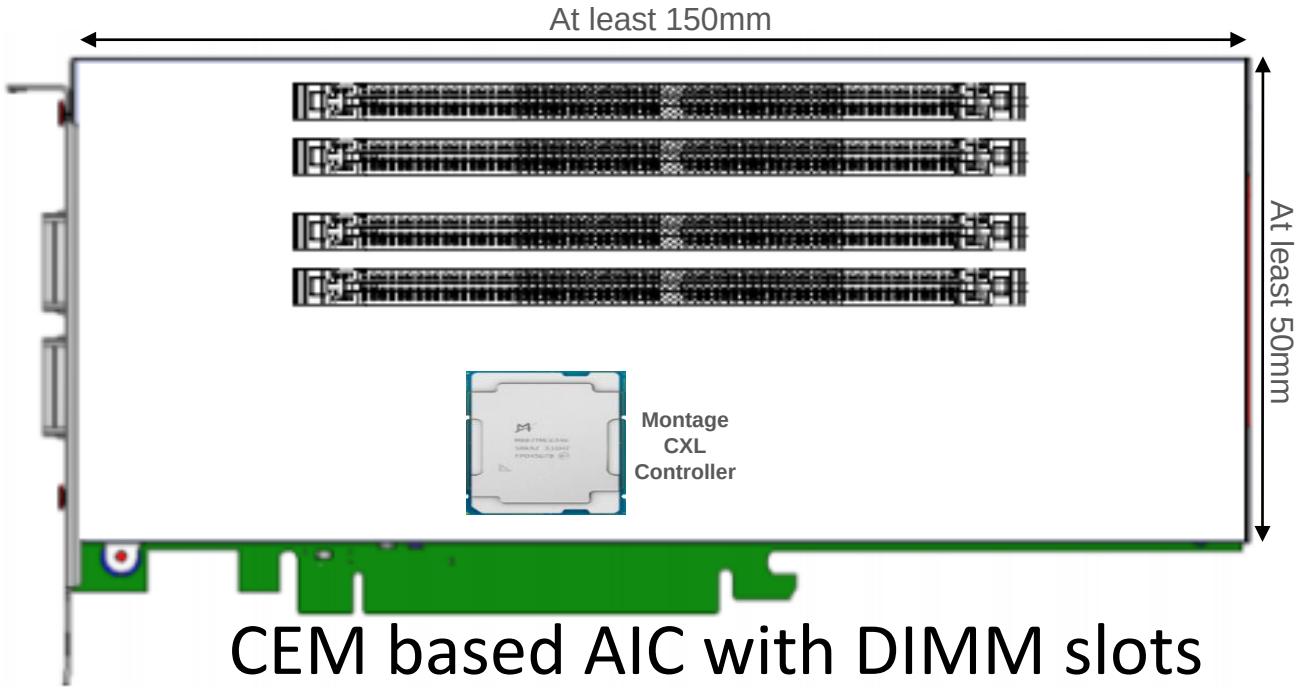

CEM based AIC with DIMM slots

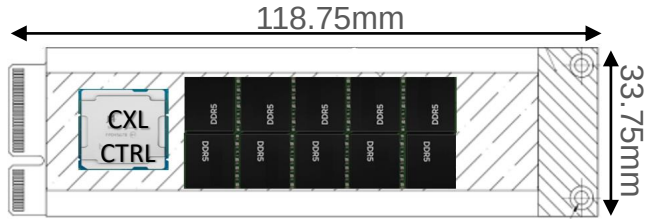

Thickness 6 to 25mm

SNIA EDSFF E1.S

AIC w/ DIMMS

Area / Size

Varied

- Comes in almost any size you want but for these applications, you would be limited by the DIMM x,y,z.
- Angled or flush mount DIMM connectors could be used to reduce Z but impacts thermals.

Montage Products:
Little MXC supports x8 PCIe Gen5 w/x80 DDR5
Big MXC* supports x16 PCIe Gen5 w/x160 DDR5

Capacity

256GB

- Dependent on # of channels & DIMM per Channel supported. Trends to larger controller.
- Max DDR speed will be impacted with 2DPC DIMMs
- DDR5 Capacity w/4 2Rx4 DIMMs w/16Gb Mono Die.

Misc

x16 PCIE
x160 DDR

- FF supports up to x16 CXL/PCIe lanes and is likely needed if 4 DIMMs used with DDR5.
- Highest Power per GB ratio – These solutions will likely be higher mainly due to PHY complexity needed to drive through a longer channel and multiple connectors.
- Volumetrically this is quite large due to the angle connectors and full high AIC size

SNIA EDSFF E3.S

Area / Size

Normal SSD

- Comes in 2 basic sizes but the largest is:
  - Width~76mm
  - Length~113mm
  - Thickness~17mm

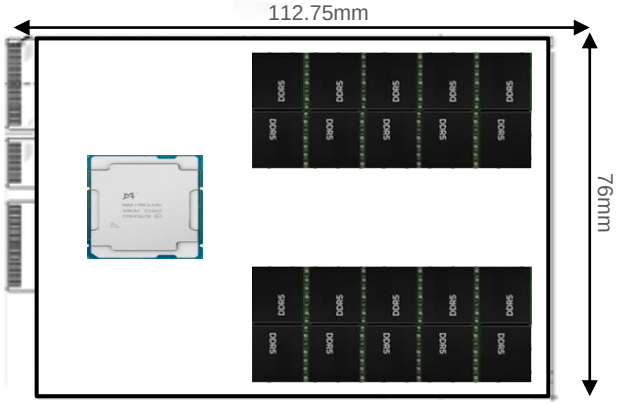

E3.S

Capacity

64GB

- Due to DDR5 package sizes, it is very difficult to place more than 20 devices on a side (40 dual sided) without more custom solutions:
  Odd channel controller support, 2nd PCB with Mez connector for more PCB area or Flex Circuit PCB.
- DDR5 Capacity w/16Gb Mono Die.

Misc

x8 PCIE

~x80 DDR

- FF supports up to x16 CXL/PCIe lanes with 160 DDR IO's though power limitations would push it into a E3.S 2T (~40W) vs E3.S 1T (~25W). This Forces tradeoff of capacity vs BW vs chassis thermals.
- Challenges persist around this form factor. Its not well optimized for traditional DRAM layouts. There is not quite enough room.

SNIA EDSFF E3.L

Area / Size

- Comes in 2 basic sizes but the largest is:
  - Width~76mm
  - Length~142mm
  - Thickness~17mm

Large

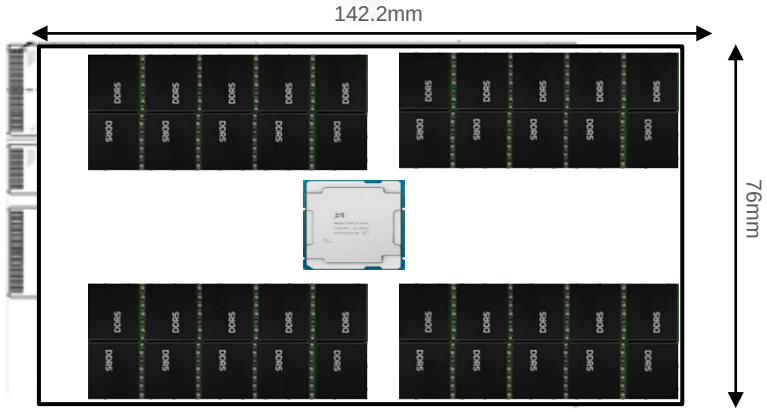

Capacity

- Moving up to the E3.L format gives just enough extra space that should be able to place 40 devices on a single side, but routing would be significantly congested. Unfortunately, E3.L adoption seems very limited.
- DDR5 Capacity w/16Gb Mono Die & 80 PKG placement

128GB

Misc

- FF supports up to x16 CXL/PCIe lanes and up to 160 IO's of DDR5 but this .L version is not as popular.
- With all the new components, routing solutions will become complicated. Pushing the controller to the middle of the package would significantly increase the PCIe / CXL channel.

x16 PCIE

x160 DDR

SNIA EDSFF E1.S

Area / Size

Small

- Comes in 5 basic sizes but the largest is:
  - Width~34mm
  - Length~119mm
  - Thickness~25mm

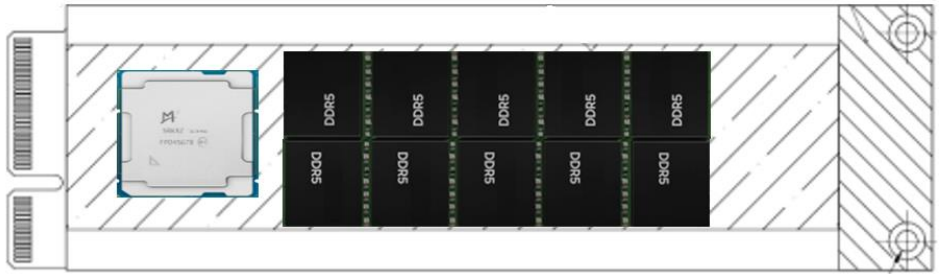

Capacity

32GB

- Due to current DDR5 package sizes, it is not currently possible to place more than 10 devices on a side (20 dual sided) without exotic solutions.
- DDR5 Capacity Max w/16Gb Mono Die @ 16 Device (ECC not counted for capacity) 3DS or higher mono density needed for greater capacities.

Misc

x4/x8 PCIE

x80 DDR

- FF supports up to x8 of CXL/PCIe lanes which is well matched for any gen DDR5 with 80 DDR IO
- Lowest Power per GB ratio – Assumption is that due to the form factor size, suppliers would likely utilize Type 4, any-layer blind and buried via PCBs allowing for very tight and optimized channels. This would allow for the best optimization of terms and PHY driver configurations.

Different routes for different solutions

Each Form Factor benefits significantly from an optimized controller layout and routing solutions.

Let’s look at those FF’s again!

E1.S Potential Layout

With the E1.S form factor space is a premium, so the package likely needs to be optimized for a Type 4. A very tight pitch and small ballmap to keep the controller area required to a minimum and make it easiest to route very congested DDR lanes…this assumes micro vias and blind and buried are all acceptable.

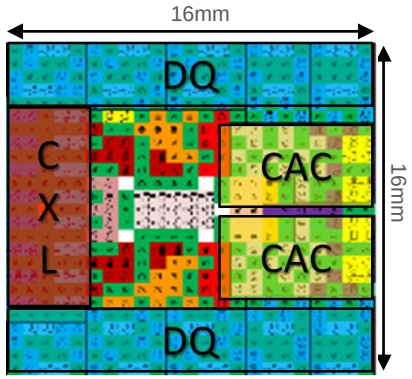

Montage CXL Controller

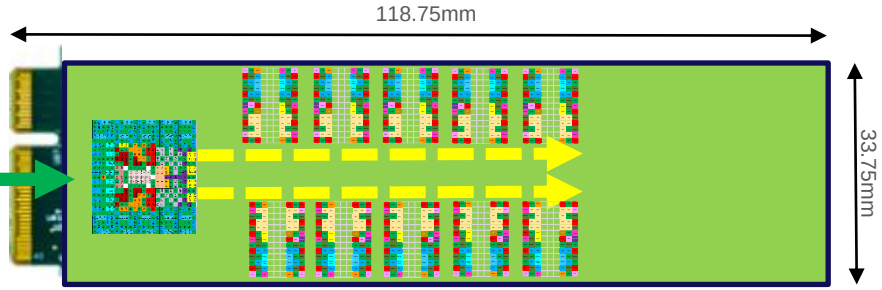

Command & Control down the middle.

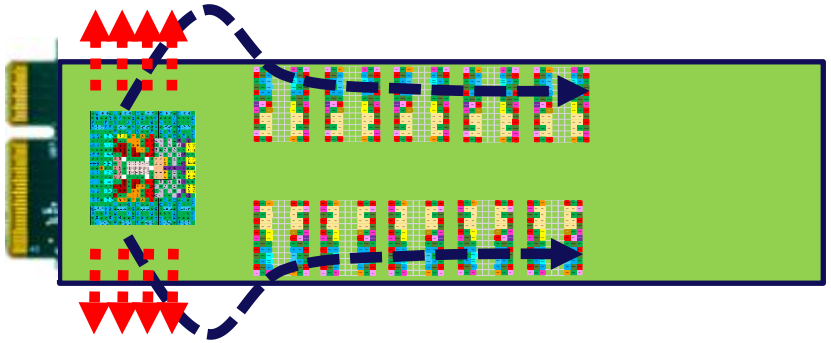

Normally the DQ's would want to break out straight in front (shown with red arrows) but DQ Channels will have to break out then dive down layers. The via field will be tight. Some vias can be in the DQ field.

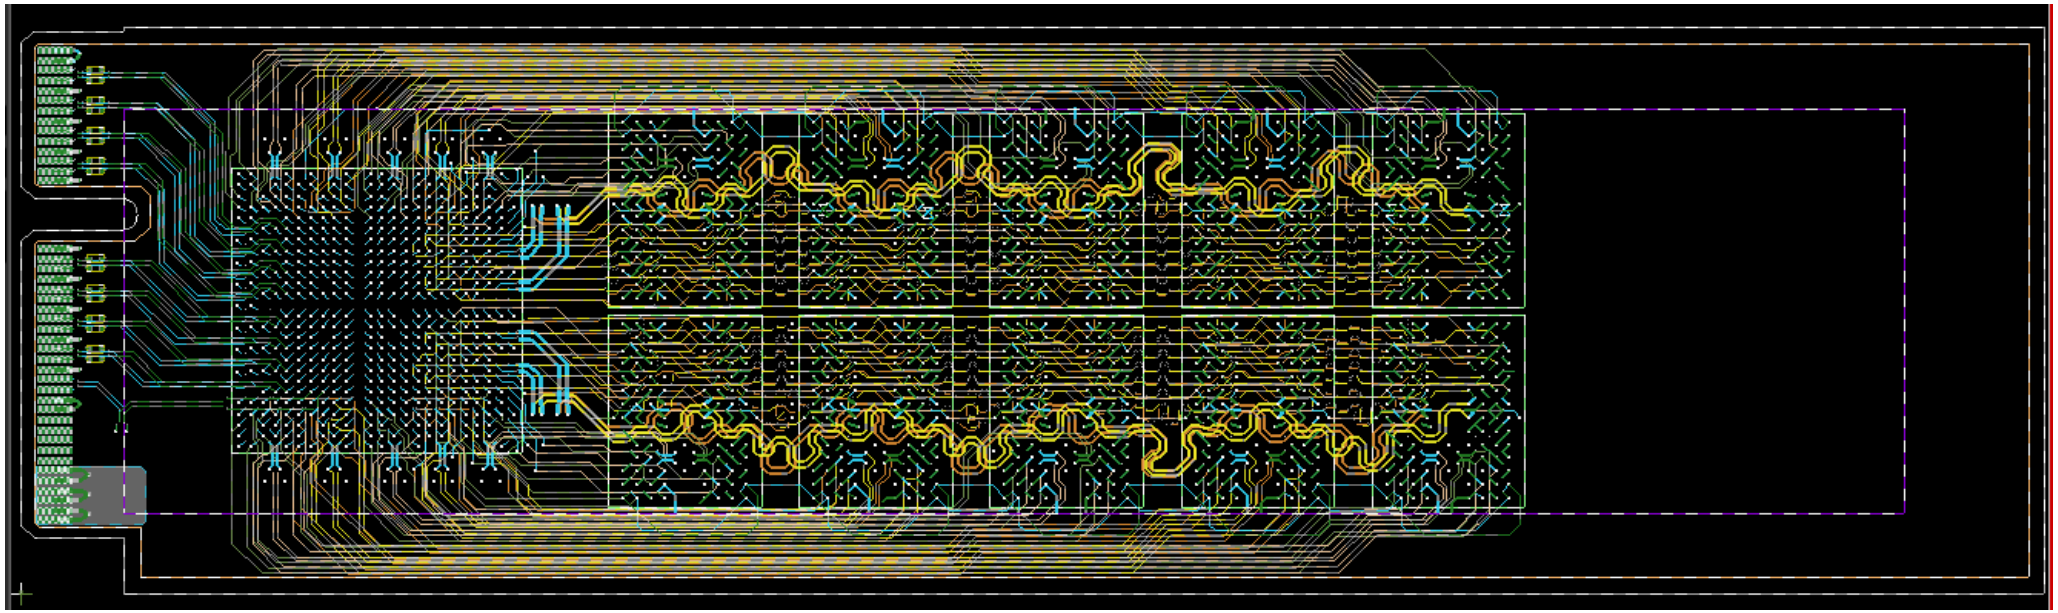

DIMM Optimized Layout

Generally speaking, DIMM solutions prefer a layout that has the DQs on the edges and the CAC through the middle. For simplicity we will call it a DDCCDD layout. This design cleanly attaches to the DIMM connector/pinout which is in the same format.

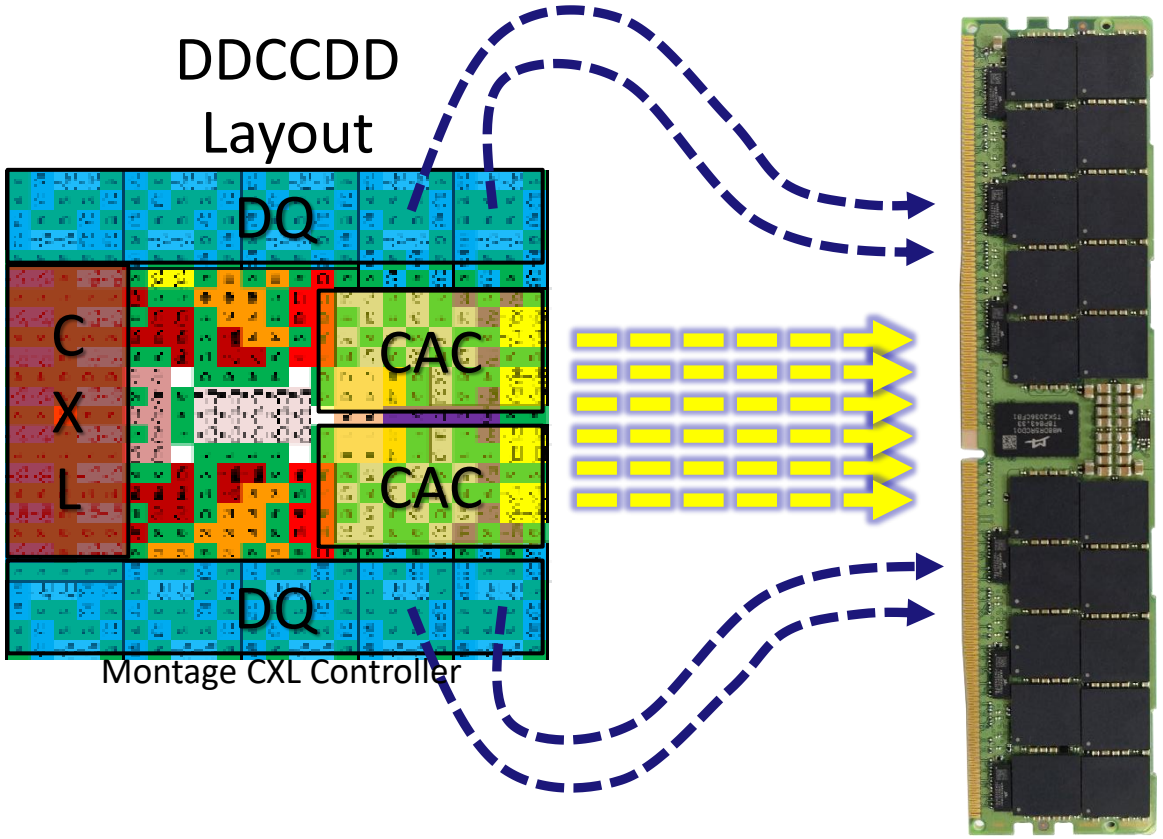

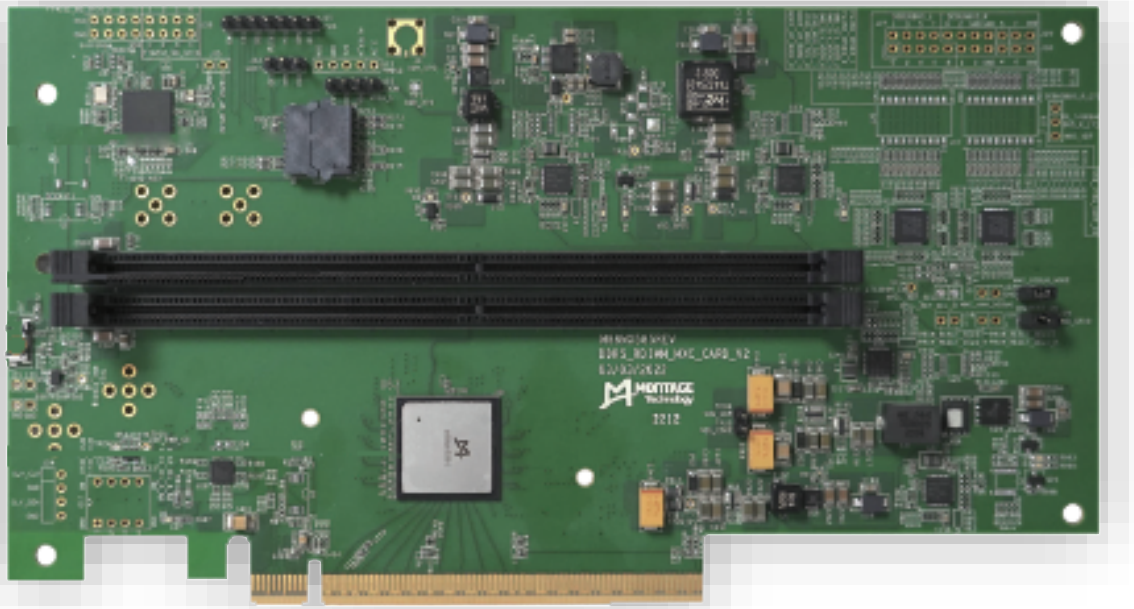

NOTE: The above image is a Montage MXC Litle Test Card & is not representative of a production card due to all the test fixtures and additional components added for evaluation usage.

What happens if we take that same controller layout and try it on an E3.S FF?

A package optimized for E1.S doesn’t necessarily work well with a E3.S

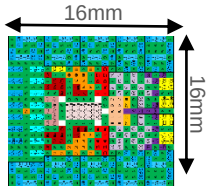

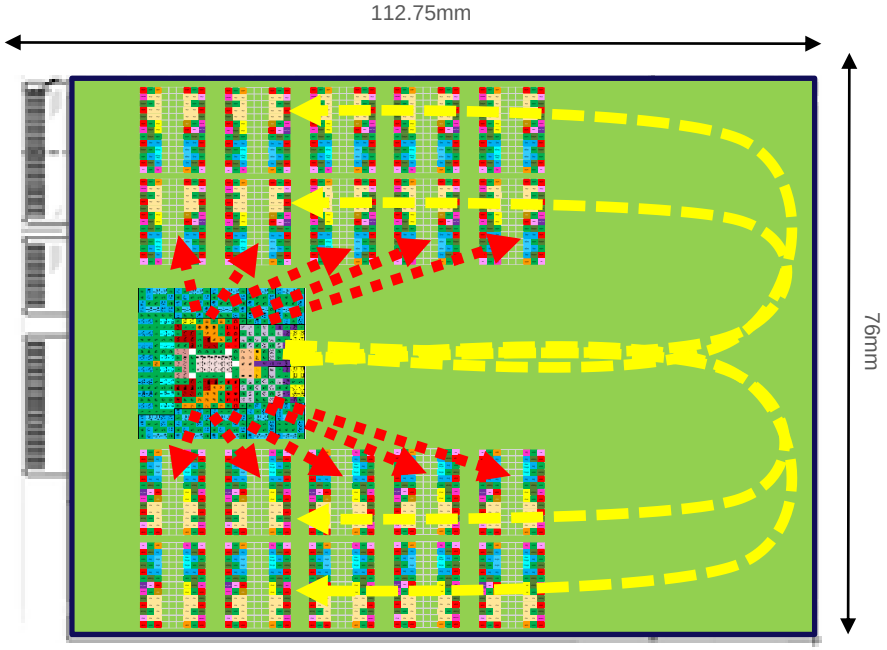

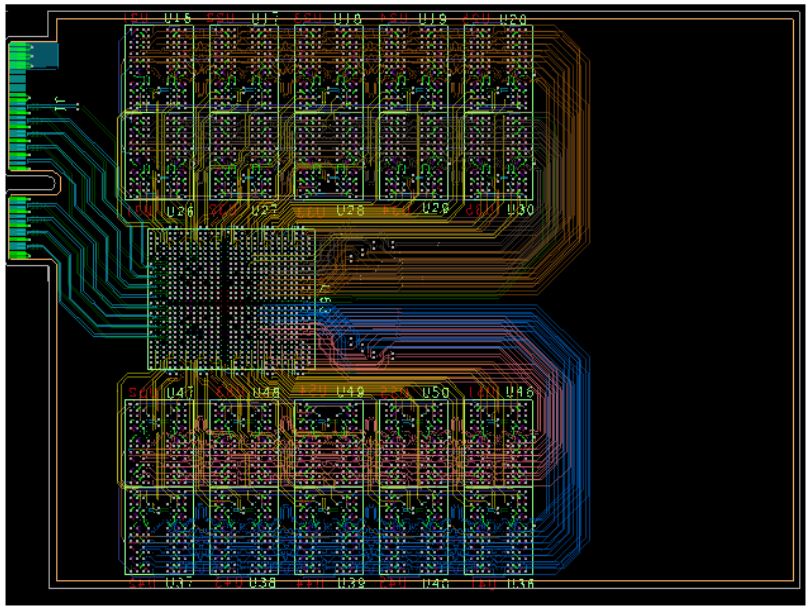

An optimized E3.S will benefit from a package designed for it

Generally speaking, DIMM solutions prefer a layout that has the DQs on the edges and the CAC through the middle. For simplicity we will call it a DDCCDD layout.

When the DRAMs are placed around the controller or in groups of channels, it's better to keep the CAC as close to the DQ groups as possible. We will call this a DCDDCD layout.

DDCCDD

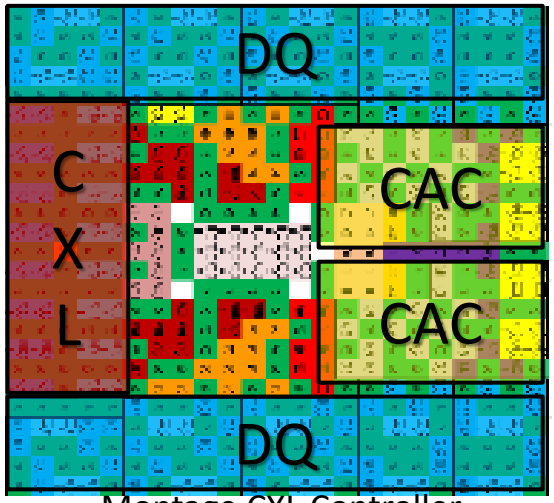

Montage CXL Controller

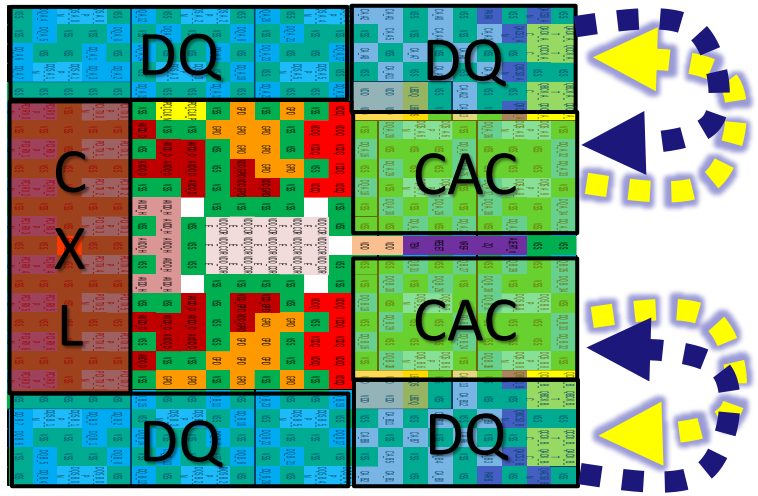

An optimized E3.S will benefit from a package designed for it

Generally speaking, DIMM solutions prefer a layout that has the DQs on the edges and the CAC through the middle. For simplicity we will call it a ODIC layout.

When the DRAMs are placed around the controller or in groups of channels, it's better to keep the CAC as close to the DQ groups as possible. We will call this a DCDDCD layout.

ODIC

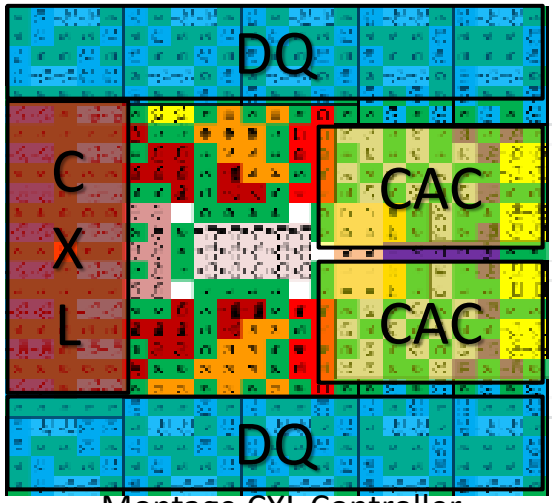

Montage CXL Controller

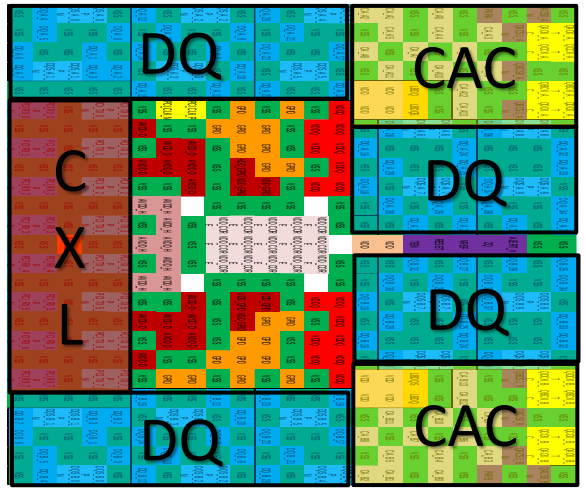

DCDDCD

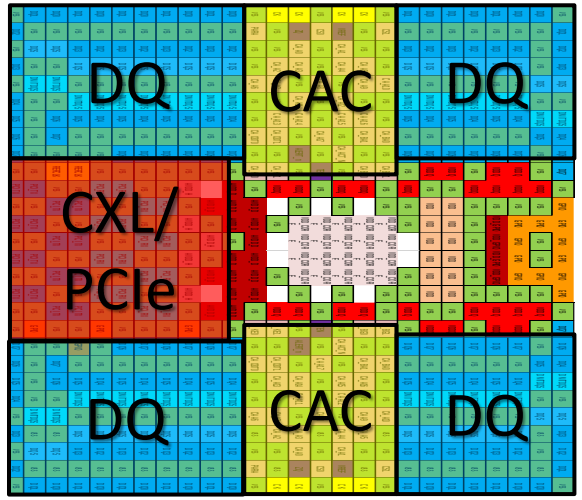

While having the CAC all grouped together in DDCCDD is ideal for DIMM slots and works well for E1.S modules, you can see that splitting up the CAC (DCDDCD layout) for each DDR5 sub-channel will make it easier to keep timings under control.
Clocks will be extremely challenging at EOL DDR5 speeds.

An optimized E3.S will benefit from a package designed for it

Since there is more room on the module, you could grow the package and increase the pitch and ballsize to make it easier to breakout of the package. Additionally, there is significant benefits of re-orienting the DQ and CA blocks.

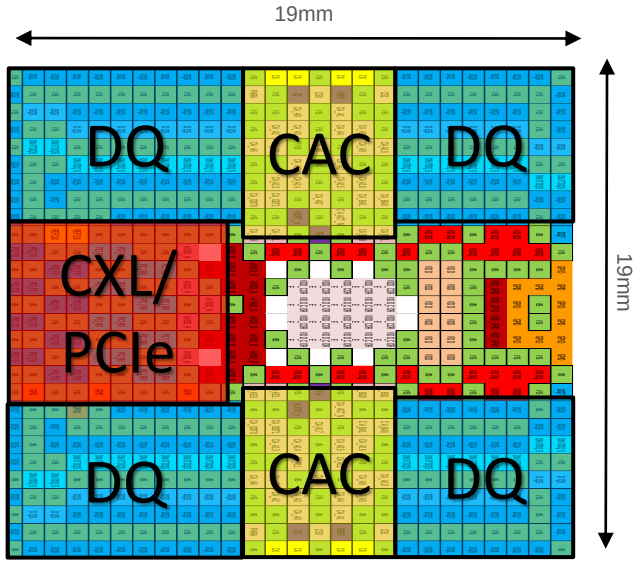

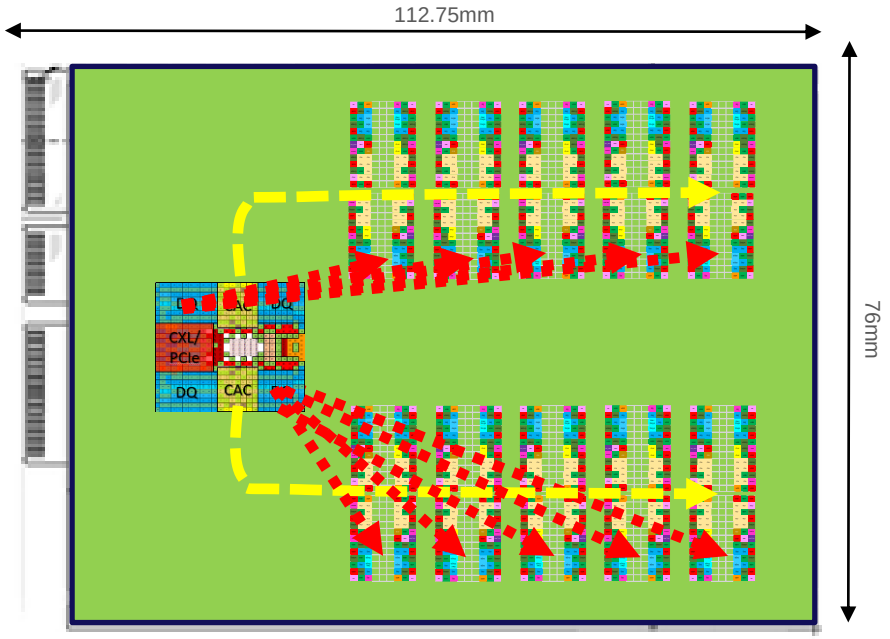

E3.S with 40 DRAM Placements on a single side

With 20 devices per side being such an easy layout, people will quickly start trying to figure out how to route 40 devices per side without usage of flexboard or other more exotic materials but from a cartoon layout, you can see how routing the CAC (CK/CS/CA) will be very messy.

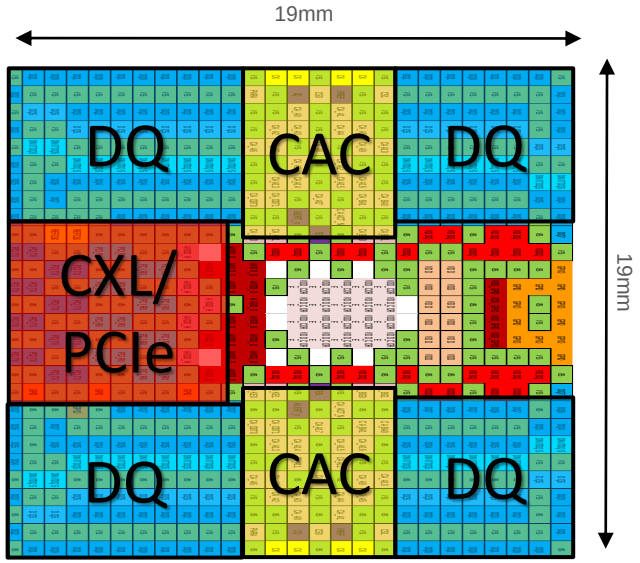

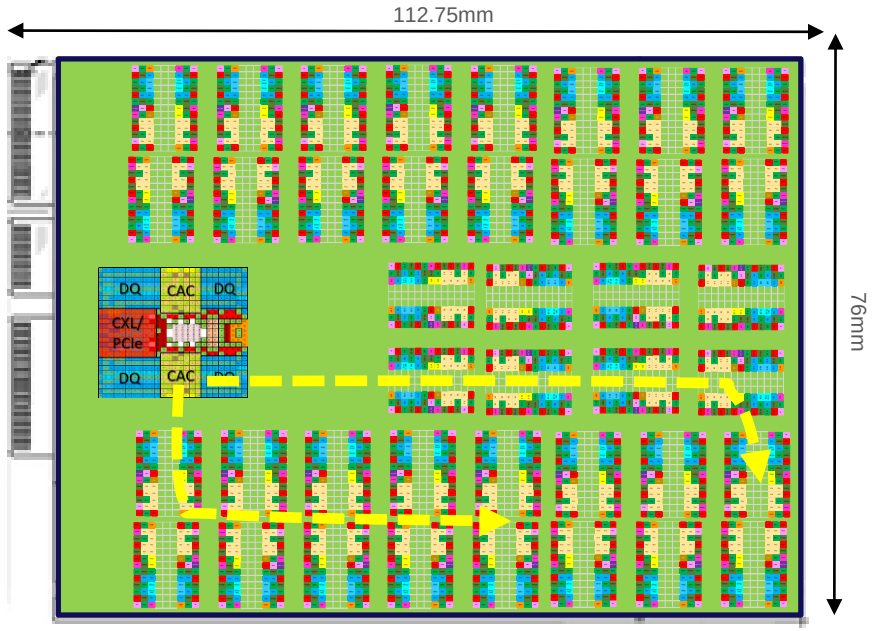

So many solutions, so many packages…

Different Package Options – Just a sample of the smaller x8 designs

Proto#1 – 511b

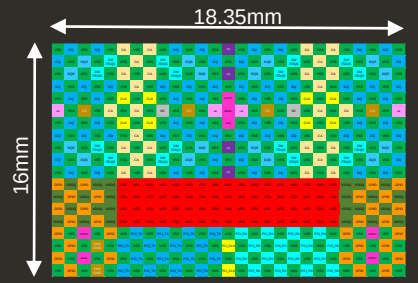

2CHx32 DDR
X8 PCIe Gen5
511 ball pkg
0.65x0.8 pitch
18.35 x 16

All Sig Gnd referenced
S2G Ratio 1:1
Power located mainly off center.
Only two different input rails shown
DDR breaks out north
Lots of grounds makes breakout much easier.
PCIe breaks out south.
GPIOs generally southern side

Extra GPIOs = 30

Most friendly for E3.S implementations

Proto#18 – 727b

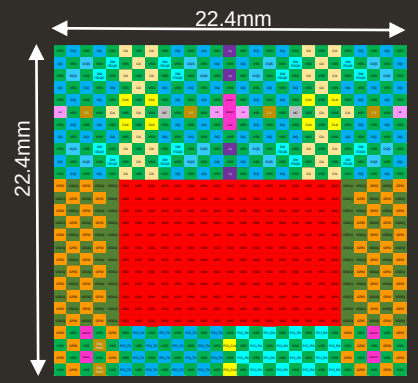

2CHx32 DDR
X8 PCIe Gen5
511 ball pkg
0.8 x0.8 pitch
22.4 x 22.4

No Signals under Die option – need actual die size to optimize.

All Sig Gnd referenced
S2G Ratio 1:1
Power located mainly off center.
Only two different input rails shown
DDR breaks out north
Lots of grounds makes breakout much easier.
PCIe breaks out south.
GPIOs generally southern side

Extra GPIOs =62

Most friendly for E3.S implementations if No Sig under die is needed

Proto#22 – 391b

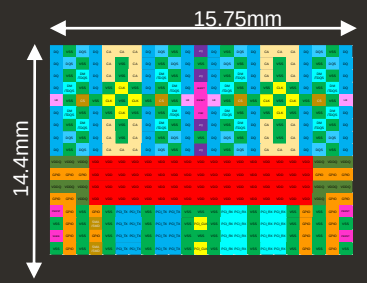

2CHx32 DDR
X8 PCIe Gen5
391 ball pkg
0.65x0.8 pitch
15.75 x 14.4

All Sig Gnd referenced
S2G Ratio 2:1
Power located mainly off center.
Only two different input rails shown
DDR breaks out north in channels
PCIe breaks out south in channels
GPIOs generally southern side

Extra GPIOs = 24

Proto#31 – 729b

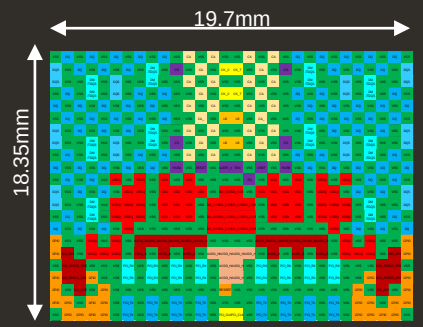

2CHx32 DDR
729 ball pkg
0.7x0.65 pitch
19.7x18.35

All Sig Gnd referenced
S2G Ratio 1:1
Extra grounds grew package
Power located mainly off center.
Several input rails added
DDR breaks out E/W in channels
PCIe breaks out south in channels
GPIOs generally southern side

Extra GPIOs = 28

Proto#36 – 378b

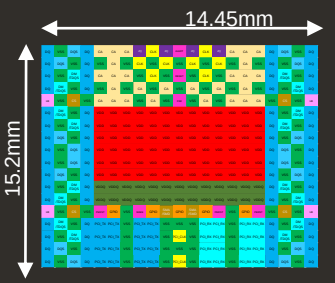

2CHx32 DDR
X8 PCIe Gen5
378 ball pkg
0.65x0.8 pitch
14.45 x 15.2

No Signals under Die option – need actual die size to optimize.

All Sig Gnd referenced
S2G Ratio 2:1
Power located mainly off center.
Only two different input rails shown
DDR breaks out north in channels
PCIe breaks out south in channels
GPIOs generally southern side

Extra GPIOs 5 – very limited, likely need to add some in power area

Proto#44 – 378b

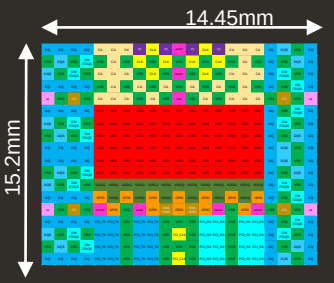

2CHx32 DDR
X8 PCIe Gen5
378 ball pkg
0.65x0.8 pitch
14.45 x 15.2

No Signals under Die option – need actual die size to optimize.

All Sig Gnd referenced
S2G Ratio 2:1
Power located mainly off center.
Only two different input rails shown
DDR breaks out E/W in channels
PCIe breaks out south in channels
Optional extra GPIOs mixed with power

Extra GPIOs = 12

Most friendly for E1.S implementations

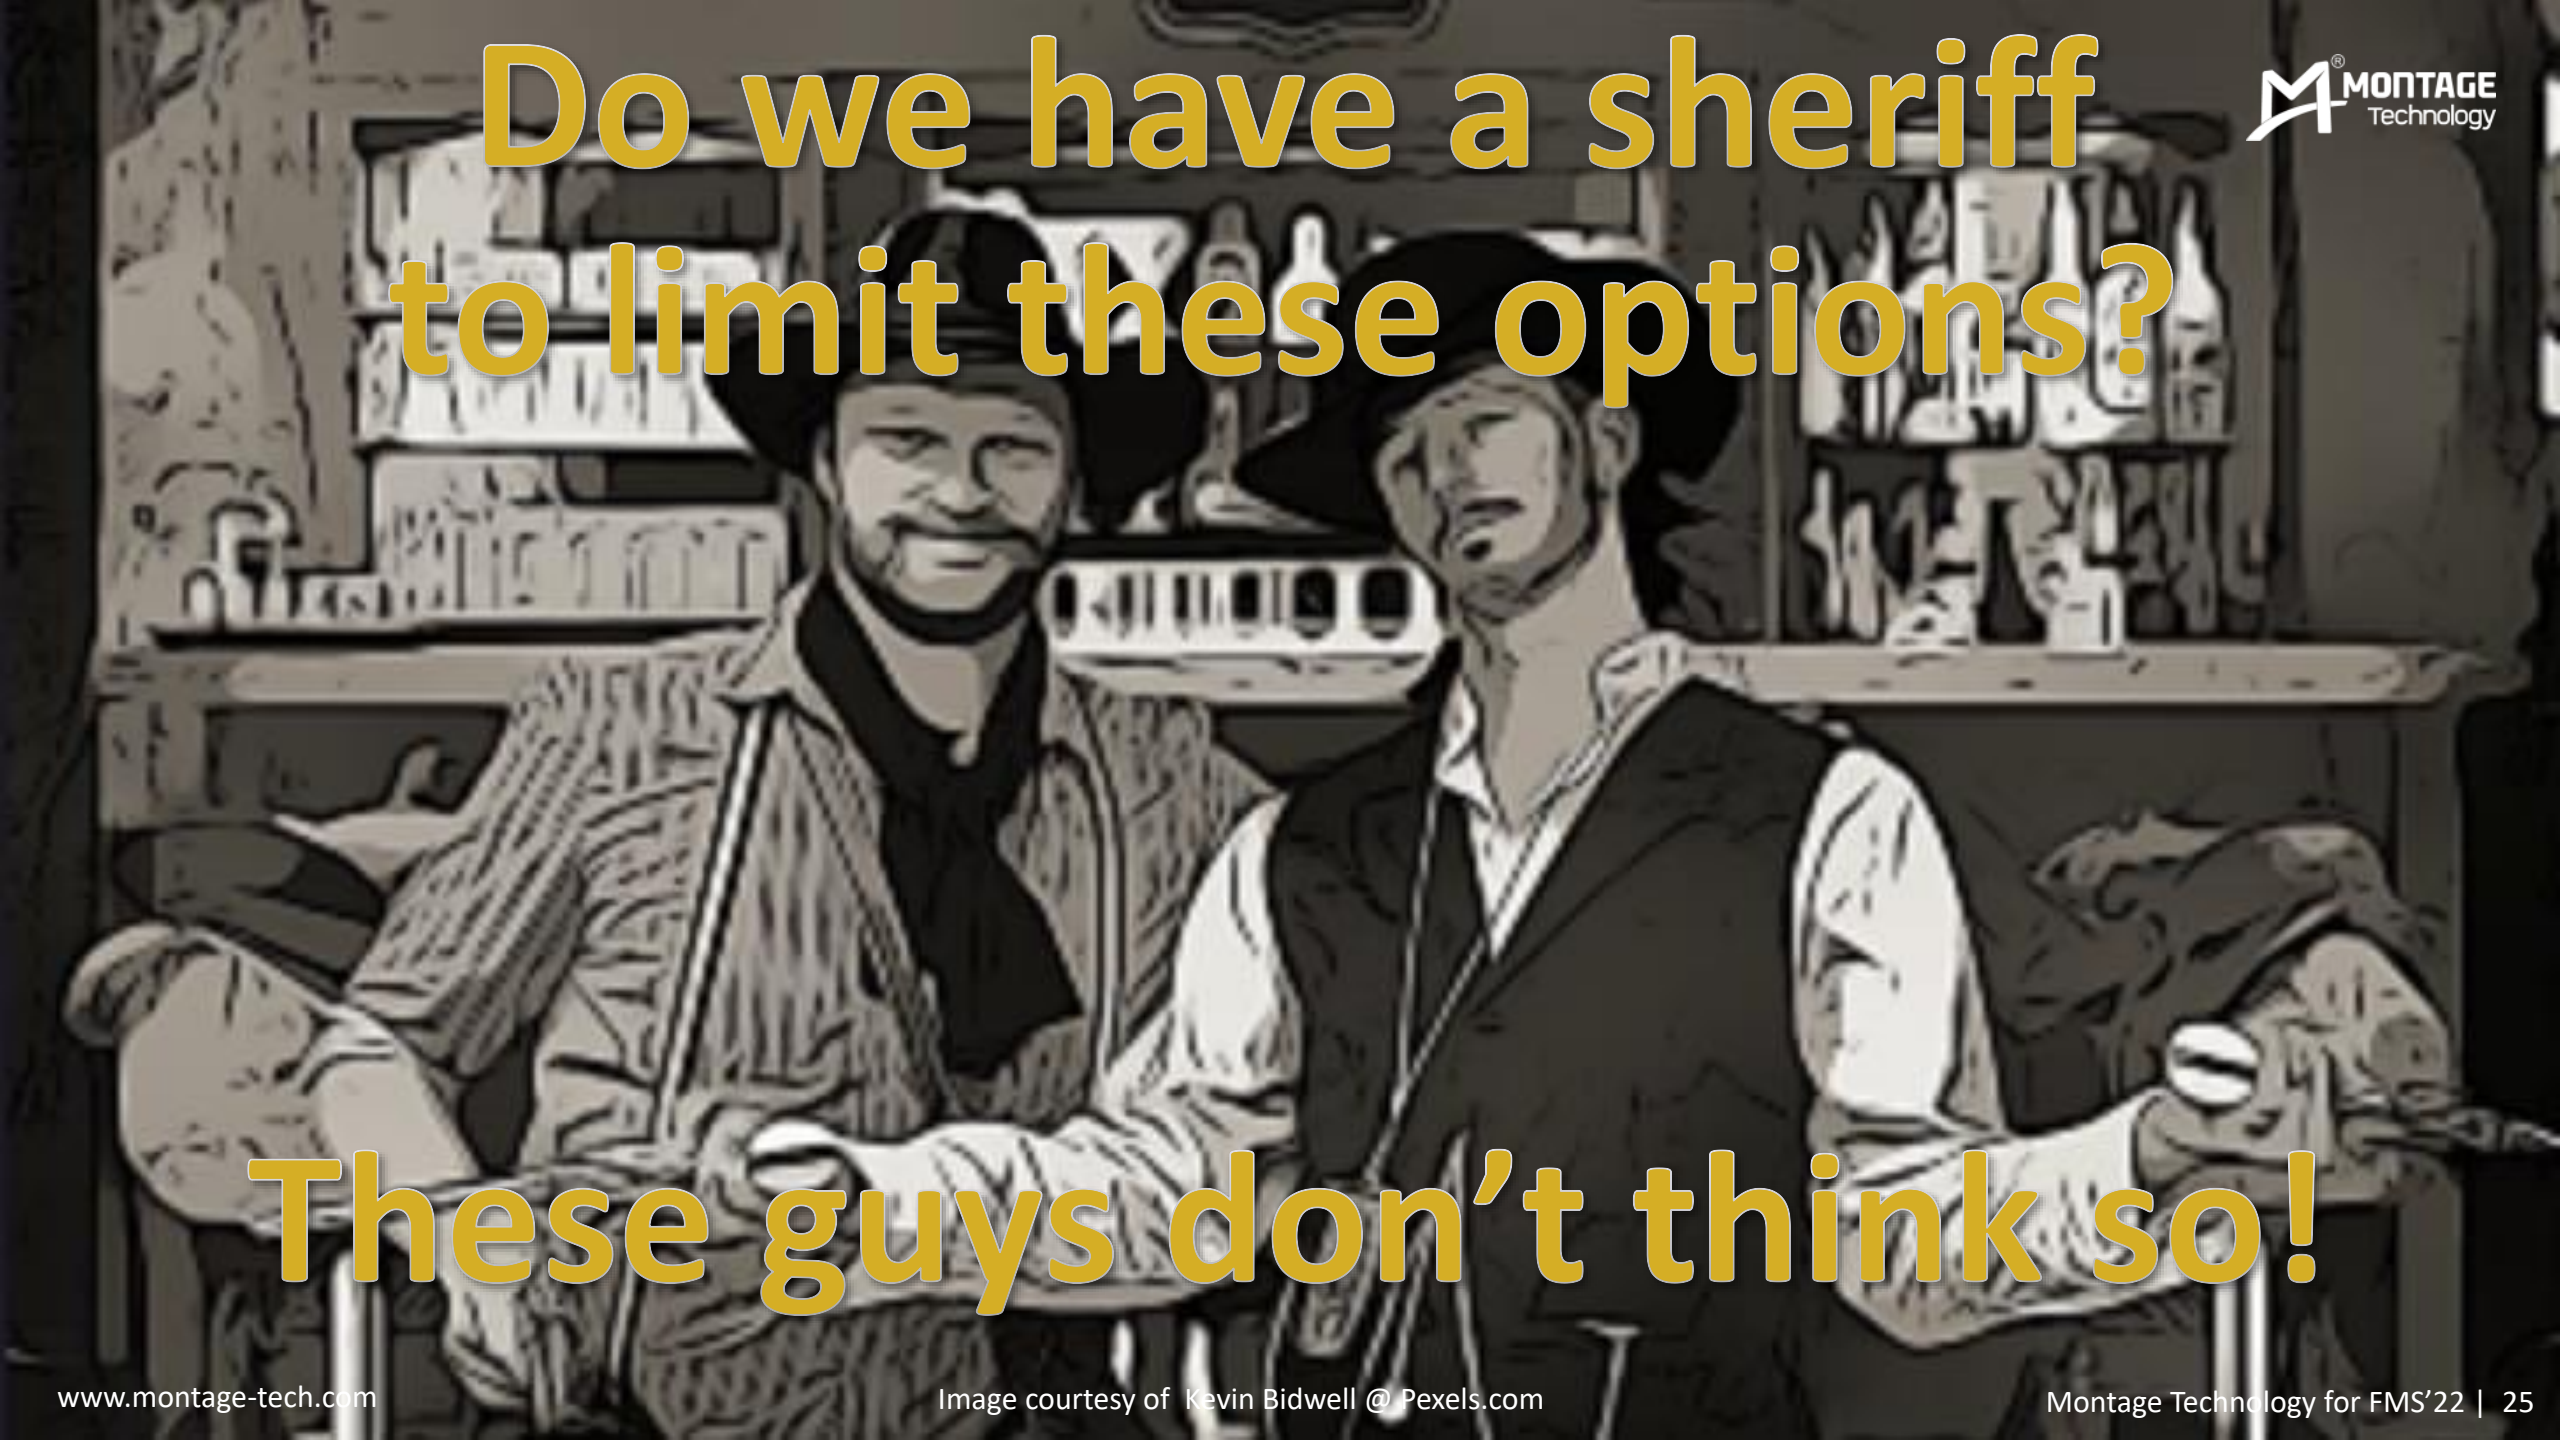

Image courtesy of Kevin Bidwell @ Pexels.com

While the CXL protocol is a fantastically agnostic to devices behind it, having so many options available makes it hard to plan for the proper controller features and layout.

Is there a form factor
style that will corral the options?

Photo courtesy of Lifesavers Wild Horse Rescue

Call to Action

There are many aspects that influence and are influenced by your form factor choice and while "*variety is the spice of life*" it also comes with increased cost and complexity to you and the industry.

There is a clear need to help the industry 'norm' around at least a couple of the form factors (and their associated controllers) to help reduce cost and make CXL DRAM Memory Expansion Solutions successful. Some obvious areas that we didn't go into details today but are still very significant:

- Thermals
- Power
- Chassis
- Performance

Photo courtesy of Lifesavers Wild Horse Rescue

THE END