

NETEVENTS

## GLOBAL PRESS & ANALYST SUMMIT

**DRAFT**

### *Conference Debate Session VIII—Finally! New Storage Paradigms for Enterprise Petabytes and Exabytes*

Chair: Steve McDowell

**Senior Analyst, Moor Insights & Strategy**

Panellists:

Kevin Deierling	Vice President of Marketing, Mellanox Technologies
Lynn Lucas	Chief Marketing Officer, Cohesity
Thad Omura	EVP of Marketing, ScaleFlux
Tom Leyden	Vice President Corporate Marketing, Excelerio
Stefaan Vervaet	Sr Dir Product Marketing & Alliances, Data Center Systems Business Unit, Western Digital Corporation

#### **Steve McDowell**

Everyone welcome. I'm Steve McDowell from Moor Insights & Strategy. Moor Insights is a boutique analyst firm based out of Austin, Texas founded by Patrick Moorhead who ARInsights calls the most quoted analyst in the world. I know some of you - many of you - have talked to Pat, right. We're here today to talk about storage.

At the keynote this morning, it was said that networking used to be really boring, but it's not any longer, right, and I think the same can be said of storage. There's a number of trends that are influencing how we deploy storage in the enterprise, and what that means to the solutions that IT practitioners are - are, you know, deploying and - and learning to work with.

The overriding trend is - is data is exploding. The volume and velocity of data arriving into the data centre is out of control, and it's coming from a variety of sources. It's

coming from the internet of things, it's coming from external data sets, it's coming from internal - you know IT and technology people, we like to measure and track everything, right. We have this explosion of data that needs to be processed, it needs to be processed close to the source with low latency and - and fast processing.

We talk about IOT. Analytics and machine learning is another disrupter that's impacting how we think about storage in the data centre, and how we think about information architecture.

Artificial intelligence is rapidly moving out of the realm of scientific computing, more into the enterprise. We're looking at all of these data, all of these data points, and enterprise is struggling with how do I apply artificial intelligence and machine learning to derive new insights.

*Harvard Business Review* put a quote out a few months ago that said one-half of one per cent of the data collected by enterprises is actually used for anything. We're seeing, you know, the AI world stood up and said well we can help solve those problems and drive new insights. Every business today is a technology business.

At the same time, you have these trends influencing what people want to do with their data, and how that influences storage architecture. From a technology provider's side, there's like a couple of trends that are also impacting and merging into how we solve these problems.

One is the rise of software defined. So we're pushing from HCI - Hyper Converged - up from the bottom, composable kinetic infrastructure from the top, and at the end of the day what we're seeing deployed in data centres for IT is really a kinetic infrastructure to use Dell EMC's term but - but one where I can leverage resources, and use software to orchestrate those solutions.

Finally, new storage technologies from media and inter-connect are really disrupting the way that we can build these systems. We - we take these new technologies, put software define on top, and it really introduces a world of flexibility into the kind of storage eco system that we have not seen before.

If you look at storage, you're familiar with kind of this data tier that we think about, right. At the top is very fast hot data that we need right now with no latency and infinite - infinite throughput. At the bottom is kind of the archival glazier, not on the chart, but at the bottom there is glazier storage that we never touch, you know, data that we keep for regulatory and other reasons.

On the media front, this has been disruptive. In 2019, the top tier we're going to see persistent memory, byte addressable, if you think about byte addressable SSD. That we live in a DIMM slot that kind of changes the way we think about server design and persistence. That impacts file systems, databases, anything that lives on your server and that's not traditional tiered storage, right, so we think about that differently.

That's been enabled by technologies like 3D Crosspoint which entails branding Optane. We're seeing that same technology today, you know, the SSD market is itself segmenting. We're seeing Optane at 3D Crosspoint being put into SSDs where I have no latency very high throughout, but it's expensive, and you don't get the densities you

get with traditional 3D NAND. At the same time the 3D NAND space is moving into the mainstream. It's fast displacing hard drives and server and storage right across the board. It is commoditising, and we are continuing to see innovations there. I think Micron this week announced their QLC technology which gives us higher densities and in greater read performance.

At the bottom, hard drives don't go away. They will always have a place in at least in the eight to 10 years because the economics are just right for that. So that is happening on the storage media front.

Interconnects are also - you know, this is probably the most exciting time since CN fibre channel arrived in the late 1990s in terms of interconnect disruption which is going to drive some of these kinetic architectures.

iSCSI fibre channel aren't going away, but they are being replaced by NVMe over Fabric, which gives me - which takes the protocol stack essentially out of the path. So I get very high performance, low latency access to external storage. NVMe over Fabric, you know, it's not just one thing. We're seeing that over InfiniBand, we're seeing that over fibre channel, we're seeing that over Ethernet. The winners - the winners will shake out over time.

Then finally, all of these technology things are working to define new storage architecture. So I put a box at the top here that says software defined, but that's inclusive of hyperconverged, composable, converged, all of these different things. What it really means is I can use software orchestration now to treat my resources as fungible to give me as an IT practitioner great leverage in how I deploy storage systems.

Those are kind of the trends that are driving disruption in the storage market. My final slide here is really a call to action for the people in this room. If we look at the Tier 1 OEMs - HPs, the Dells, IBMs, the [wallways] of the world - yes, they move markets and they drive markets, but the trends that are influencing them long term, some of those come out of those organisations, but most often they come out of them more disruptive and non-Tier 1 players. I think as press and analyst, I'm certainly guilty.

We generate a lot of ink for the big names, but I think that does a disservice a little bit to our ultimate consumers of the information that we put out, right, the IT practitioners, and that storage disruption, disruption across all the technologies, is very often not from the Tier 1s.

With that background, I'm going to let the panel introduce themselves, and say a little about who they are and what they're doing. We'll start with Kevin.

### **Kevin Deierling**

I'm Kevin Deierling, I'm the VP of Marketing at Mellanox Technologies. We're a networking company, and we're really embracing all of these new faster storage. We see that hard disc drives had 10 millisecond access today. These new technologies like NVMe over storage, and also now the next generation of persistent memory is about 10 microseconds, so that's a thousand [x'd] improvement in performance over the last

decade. That's the difference between driving from Boston to here, and driving up to Levi's Stadium in terms of the change in the amount that it takes.

What that means is faster storage needs faster networks, and that's really where we come into play. We see the fibre channel continue to decline about 10 per cent a year, and we see a new Ethernet storage fabric emerging, which is a converged fabric hyperconverged.

There is no fibre channel in the cloud, and now we see that happening in the enterprise in Tier 2. So we're excited to deliver that Ethernet storage fabric with partners like HPE that have M series switches based on our technology.

### **Lynn Lucas**

Hi I'm Lynn Lucas and I'm CMO at Cohesity. I feel familiar with Nutanix. Our CEO, Mohit Aron, was the CTO and co-founder of Nutanix, called in some cases the father of hyper-convergence. He observed in the secondary storage market that it was time for a radical transformation.

Cohesity is creating the category of hyper-converged secondary storage, consolidating all of the silos for backup files and objects, test dev and analytics on one hyper-converged platform, and then allowing businesses to really do a lot more with their data and their applications, moving them seamlessly across their private data centre to the public cloud, and having a lot of - more access to what's in that data, especially as we see the rise of machine learning and - as we talked about yesterday - GDPR and the need to really understand all of the information that you have.

### **Thad Omura**

I'm Thad Omura, I'm the EVP of Marketing for ScaleFlux. ScaleFlux is the pioneer in deploying computational storage at scale. You may be asking what's computation storage?

Computational storage is essentially a new paradigm in storage where we're adding compute functions right into the storage drive itself. We're doing that because we see the rise of data-driven applications. Applications like - that run on database platforms, big data, content delivery, AI and ML.

They are all experiencing massive data growth as we've been talking about this whole conference. But what you want to now do is derive insights from that data at a much faster pace. We're now seeing that as storage medias are transferring from hard disc drive base to flash, and the connectivity to those drives has increased in performance, moving to PC Express and NVMe. It's now driven by the applications and the technology to be able to compute locally the data, and not take all of the time to move the date to compute. So ScaleFlux again driving compute to data on - on flash storage technologies.

**Tom Leyden**

My name is Tom Leyden. I'm in charge of marketing at a company called Excelero. What we do is shared accelerated storage. We enable companies to leverage very fast storage devices at - in a distributed way, specifically NVMe.

You all know the benefits of NVMe, very low latency, high IOPS. When it's used locally, our software solution allows customers to deploy NVMe in a distributed fashion. It can be in just a bunch of flash or can be in the segregated way as well. Customers will get the performance, the latency of the NVMe as if it were used locally. We only add about seven microseconds of latency when using NVMe over the network. Now we see that as a great solution for several use cases. A few verticals where we are seeing great adoption is media and entertainment, HPC, but also in AI machine learning more and more.

**Stefaan Vervae**

Good morning my name is Stefaan Vervae based out of San Jose. I'm Western Digital, a small disk manufacturer. We are transitioning in Western Digital - I do not know if you guys have seen it, in the last three years.

We are very well known as a disk manufacturer, the largest one after the acquisition of San Disk, we've also adopted Flash. Today I work for the DCS, the Data Centre Systems Business Unit, and I run Marketing and Solutions and Alliances. Our goal in the last three years has been to move up the stack, and kind of bring better solutions to our customers. So what does that mean?

Today we have actually built a new business unit that actually delivers scale-out systems to our customers, specifically customers that are looking for a fully integrated platform. That could be NVMe based, disk based, to a fully integrated object storage with a restful API. We are focused on large and structured data sets, specifically HBC type of use cases, AI machine learning.

As you know, we sell pretty much everything that's under the sun today. We sell from embedded devices all the way down to high enterprise products. The DCS, the Data Centre Systems Business Unit, was kind of created to bring enterprise solutions to market, that kind of bring those - those - all these innovations together from SMR, PMR drives, to full NVMe over Fabric in a platform as a - even a storage server, so we're actually moving aggressively.

All my colleagues here at the table here, we are collaborating in some form or fashion with the breadth of eco system players that you see out there, whether they're software defined players, or they're hardware players.

**Steve McDowell**

Awesome, thank you guys. Hey so as we look across the board right, there's a lot of disaggregation in the models that you all are - are providing. We talk about NVMe, which takes latency out of the storage path in the box. We talk about NVMe over Fabric which takes latency out of the data centre storage path, and we talk about software

defined, which gives you the flexibility to deploy storage solutions in a way that fits the workload, whatever that workload might be.

My question to the panel is is direct attached storage dead? Are there new models emerging that replace that or does it continue to have a place?

**Kevin Deierling**

So I'll start. Direct attached storage is not dead it's just transformed. The real issue is you can't have a box where you have really expensive NVMe over flash drives that's running at 30 per cent capacity, and another server node that's at 100 per cent and needs 50 per cent more. It's sitting as a stranded asset in another server. That's really what this is all about - is having a shared storage resource that you do not care and the hyperscalers realise this. When they're solving big data problems, the data set is larger. They kept building large and larger super computers. That's what Mellanox also does, that scale up, and at some point, the room was not big enough.

You couldn't build a bigger super computer to contain all the data that they're processing so they scaled out. Today you don't care whether the data is on your server direct attached, or whether it's at the other end of a football sized data centre. You need to be able to go grab that data, and process it, and take it from hundreds of different nodes.

Direct attached storage is still there but it has to be shared, and fibre channel isn't going to cut it. It's, you know, we deliver at Ethernet three times the performance for one third the price versus fibre channel that has to be converged and hyperconverged and that's just the way the world has changed.

**Steve McDowell**

So the fabric of the new data centre is Ethernet?

**Kevin Deierling**

Absolutely.

**Steve McDowell**

Okay, go ahead.

**Tom Leyden**

Yeah, I agree with what Kevin said. What we see in the world today, how companies can be competitive or more competitive than others is actually by having information, by having an ability to process a lot of data, so the leaders are those companies that are able to process more data and process it faster. To do that they need more scalable and faster storage.

Now the faster storage fortunately - thanks to friends, for example, from Western Digital - comes in the form of NVMe, but you also need new architectures. What's important in those architectures is that you work away bottlenecks. A very popular bottleneck - well not popular but - but a very common bottleneck in architectures that

we've seen until now - are the storage controllers. What we need actually is a form of shared accelerated storage design and architecture where there's no controller, and where these fast storage media are accessible across the network, across fast networks preferably.

### **Lynn Lucas**

I would just add, I think, speaking to the concept of software defined. Customers don't care, I think as the gentleman mentioned, really where at the end of the day the data is, and so it's the job of these new architectures to make that seamless for them.

We have an analogy of, you know - especially in this area of the cloud of, you know you - the data centre can be thought of as your home, and the cloud can be thought of as your hotel, and the world's a better place because you have both.

Sometimes I want to go visit a hotel but I might not want to live there permanently due to the expense, but it's nice to visit. Having a software-defined model really allows the IT organisation to start to abstract that, and provide to their business clients the services and the data that they need independent of them really starting to worry about where that data is, and I think that's the real benefit of what software defined is bringing.

### **Steve McDowell**

So Lynn and Kevin, you both - you both alluded to customers don't care where my data is, but I think IT does care. When you look at workloads - and software define is increasingly driven by, I think, the server side of the house less so that the storage side of the house which clouds their thinking.

We are used to a world as server guys thinking about, you know, VM ware and V motion, and I can move processors wherever they make the most sense, and create a dynamic infrastructure that way. Very often, there's this long tail of gigabytes, or petabytes, or terabytes of data that have to flow with that. In the world that we're talking about, that data may be at the edge, it may be at the cloud. If it's in the public cloud maybe it's on Amazon East, or West, or Asia, and latency is what defines success and analytics for sure.

How do we make sure that the data is where it needs to be relative to the processing?

### **Lynn Lucas**

Yeah, a couple of things that you brought up there is I think you are absolutely right. IT absolutely cares where it is. I think that business owners should not care, and that's the vision. They care because I think many ran to the cloud and found out how expensive that could be, and the difficulty of both getting it there and getting it back in a timely manner.

Now we - you touched on two concepts. We - we define secondary as non-mission critical, non-real time. Just to be clear, we are not talking about like a real time analytics application, or real time airline ticketing, or things like that. I think to your point, data's being generated everywhere, right. You know, your car's becoming a mobile data

centre in effect, and with these very large data sets, as you point out, one of the philosophies has got to be that you can't bring data to the compute, you have to bring compute to the data, and that's another reason for software define model, because it simply won't be tenable in some environments to do that anymore.

### **Thad Omura**

I think that ultimately it's the application requirements that drive that location of the data, and for being a little bit controversial we do still see a large population of folks who want to continue to have more hyperconverged platforms for the ease of manageability, and there is a transition for some of that to go to disaggregated model, but we're in a transitional period. It really depends upon the IT shop's capability to manage a new type of infrastructure versus just buying the same box and replicating. There are new technologies to treat hyperconverged architectures as flexibly as disaggregated models.

When we go back to where the storage is to - what ScaleFlux is seeing is it does matter because it depends upon the latency at which you can process that data. If there's a competitive advantage to be able to process within microseconds, you're going to have that data as close to the CPU as possible, or you're going to look for architectures to - as we are all talking about - move compute even closer to the data which you would do with computational storage. If things can take microseconds to even seconds, now you start to look at the storage can be located potentially farther away, then it's less disruptive to the competitiveness of that particular IT shop's and where they have that data located.

### **Stefaan Vervaet**

Yeah and I think you also see - what we're seeing on the Western Digital side is a change in software defined storage systems, right. If you look at five years ago, I used to work for a software defined storage company myself with my dear colleague Tom here.

Today where initially it started a lot in the complex stuff consolidating storage, you know tape replacement, multiple second tier third tier storage consolidation. Today it's very much driven by the application re-tooling that's happening as well. We actually have customers that come to us and say look I love your object store, we've consolidated a lot our tiers. The second challenge we have now how do we use it, and then they go down this path of re-tooling all their applications with containers, micro-services, which is fantastic. Typically driven by the need to be able to scale their compute, their new applications, their new business models that they are trying to build that are typically driven by machine learning initiatives.

Typically in that case we are - we are coming with our NVMe drives because we - they want to accelerate their - their data - their access to data, they want to bring it closer to the CPU, but to Tom's point is that's where you have to have these new types of software solutions that actually disaggregating give you the ability to access that data much faster. So the replacement of the traditional SAN as we know it today is definitely happening.

**Kevin Deierling**

Yeah I think that the only IT guys that really care about where the - the storage is is the fibre channel tree huggers, and I think ultimately they're going away because what you really need to happen is see that convergence in the network occur. We see it with hyperconverged infrastructure, for example, where you're combining the compute and the virtualisation, the storage, all in a single box, and then you're networking that with Ethernet.

So over time what we're seeing is that the software defined storage mechanism, you know, hyperconvergence makes the infrastructure go away, makes it invisible. What we're going is making the network invisible too because ultimately people care about business processes and making sure that things are working not where their data is.

**Steve McDowell**

So we talk about NVMe taking the latency out, we talk about there's certain classes of applications - and Thad I think you alluded to this - which require as little latency as possible, if I'm going financial tech or some other types of things. I think ScaleFlux and Excelero you both - you're both working to bring extremely fast storage to the processing workloads that require that.

I will ask a controversial question, 2019, 2020 - does persistent memory in - in putting my persistence in a DIMM slot solve that problem for me?

**Thad Omura**

From ScaleFlux's perspective, actually the advent of persistent memory only exacerbates the problem in terms of having more compute element actually closer to the data, because what you're now doing is actually even speeding up the accessibility to even more data, at a faster pace. So if you want the platform to scale you still need to think of architectures to move now, compute even closer to - into that persistent memory. One thing that ScaleFlux is doing is remaining extremely flexible on its media attach, if you will, to our computational storage.

Today is obviously predominantly - we use 3D NAND flash - but certainly looking forward, you would see either hybrid architectures with persistent memory NAND flash, or potentially offer system memory.

**Tom Leyden**

I agree and it goes back to what I said earlier. It is about a capability to use faster storage, and also to be able to scale. So what I failed to mention on what Excelero enables its customers to do, is we enable them to - to share NVMe across the network, but not just NVMe. We are independent of hardware and I'm trying to not use the software-defined word, but I guess I just did. Really what we offer our customers is the ability to - to use the fastest state of the ark storage and networking components to build the storage infrastructure that delivers the performance that their applications need, but with standard network - standard server components.

Essentially moving forward, if they want to use storage class memory because it gives even better performance lower latency, they will be able to do that, and they will have the right storage architecture to do that. So, yes there is a future for that. Fortunately, right now the - the capabilities of - of NVMe are mind-blowing, and I think we are good focusing just on that for the near future. We want to focus on where the opportunities are today, and not where they are going to be five years from now, even though we are ready to deal with those when they come.

### **Kevin Deierling**

I think with persistent memory, what you are really doing is changing the granularity of the data that you're accessing. If you think at a big level you access a four gigabyte movie, okay, and then if you're doing a block level access its maybe 512 bytes or 4k bytes. At persistent memory now, you're talking about byte access. What really matters there is the overhead to actually access, if you're doing really tiny amounts of data what's the overhead.

One of the downsides of software defined storage is that you're using a really, really expensive storage controller which is typically an Intel x86 processor to actually process that. That can work fine when you're grabbing a giant file - say you're downloading Netflix or something - but when you're talking about persistent memory and doing really small amounts of data, the overhead needs to go away. We've introduced things that are smart necks that actually are intelligence storage controllers so that you can hide all of that latency, and you're actually able to take advantage of things like persistent memory, so that's the real key.

### **Steve McDowell**

We are close to out of time. I wish we had more than 30 minutes but let's open this up if there is a question from the audience, something that I didn't ask.

## *Audience Q&A*

### **Audience member**

[Unclear]. Actually, I was wondering, I've listening carefully with attention and I never heard the word Blockchain. I ask this because if Blockchain really takes off, quite a lot of the things that have been discussed could be really, fundamentally changing. The streams, the quantities, hot and cold will be completely overturned. Computational needs near the data seems to be like no brainer, so I was listening and never heard the word Blockchain, and it seems to me just are you kind of thinking no that won't be, but if it really takes off how will you be reacting?

**Stefaan Vervaet**

That is an interesting one. We are doing a lot of - definitely Western Digital - Blockchain is technology we are tracking absolutely on different levels. We're actually doing some research on that which I can follow up with you afterwards.

In general we see it very prevalent in some of the use cases, like in live sciences, or healthcare, verticals where data has to be shared across different departments. We work with technologies today, and some of the start-ups are embedding Blockchain technologies as a replacement of their traditional Merkle Tree databases as a way to have built-in audit capabilities, and tracking capabilities. Some of the typical use cases there are government institutions even within, you know, within a government you have different agencies that don't trust each other that have to collaborate.

So there's different use cases that are driving faster the Blockchain technologies in the software stacks that are running on some of the storage technologies that we provide. Specifically in those use cases, they are doing some accelerations where GP use to accelerate that effort but from - from our perspective we haven't - so we're still working on what does that mean to Western Digital, what does that mean to the different memory class, and how can we contribute to that. We are doing some effort on that which I can take back and take that off line with you. We haven't really taken an approach so far in our solutions to market that would specifically say this is a solution for Blockchain, or this is how we - so that's still in the works, yeah.

**Audience member**

I am actually glad that you bring it up. Two years ago when I was interviewing with our president of the board he - he was asking me about the use cases of software for the accelerate technology, and I gave him the use cases, supercomputing et cetera. He said no but - the future - what's the future, what are the future use cases that you see. Fortunately I had just been reading up on Blockchain, so that was the one that I gave him. I think our architecture is perfect for Blockchain, especially as we are going to be able to - to deploy over multiple data centres. But like I said earlier today, as a start-up company that needs to business, I want to focus on use cases that happen right now and that is why you've heard me talking about media and entertainment and supercomputing and even artificial intelligence and machine learning that's what's happening today but trust me we're ready for Blockchain.

**Steve McDowell**

All right. Is there any more questions from the audience? Okay.

I want to thank the panel for coming out. I hope what all of you take away from this is that storage is interesting for the first time in a long time. There is a lot of disruption being driven by all of the things that we talked about. So pay attention to storage and pay attention to information architecture in IT.

Thank you guys and I think we will all be around during lunch if you want to chat more.

[end]