S1 EP16 - The Algebra of Storage Fabrics

Friday, May 27, 2022 · 20:58

Nishant Lodha, Director of Product Marketing – Emerging Technologies and Brian Beeler, Editor in Chief at Storage Review discuss the Algebra of Storage Fabrics. Brian is one of the most renowned voices in enterprise storage and the force behind StorageReview.com, a world-leading independent storage authority, providing in-depth news coverage, hands-on evaluation, detailed reviews and consulting on everything enterprise storage - arrays, hard drives, SSDs, networking and storage fabrics. Join in on their conversation discussing some of the latest trends in NVMe-oF, Fibre Channel and much more! Read the Storage Review Article Marvell Doubles Down On FC-NVMe: https://bit.ly/3KKo8C6. Learn more: https://bit.ly/3N4JRpz

Speaker

Nishant Lodha
Director of Product Marketing – Emerging Technologies

Brian Beeler
Editor in Chief at Storage Review

Host

Christopher Banuelos
Senior Manager of Global Social Media Marketing

Christopher Banuelos 00:04
Welcome to the Marvell Essential technology podcast. I'm your host, Chris Banuelos. Today listening in to a conversation with Nishant Lhota, Director of Product Marketing at Emerging Technologies and Brian Beeler, Editor in Chief at storage review. Brian is one of the most renowned voices in enterprise storage and the force behind storagereview.com, a world leading independent storage authority providing in depth news coverage, hands on evaluation, detailed reviews, and consulting on everything enterprise storage since 2010. Be sure to subscribe to our podcast to stay up to date on future episodes.

Nishant Lodha 00:38
Hey, Brian, welcome. And it's so good to have you on the Marvell Essential technology podcast. Where does this podcast find you?

Brian Beeler 00:46
Well, I'm here in Cincinnati, but I can tell by the sun coming through your window that you must be somewhere much more springy than I am

Nishant Lodha 00:54
Sunny California. It's been raining for last couple of days, which is which is really good.

Brian Beeler 00:59
The one thing though, that I know I'm excited about and we haven't talked about this much yet is seeing people in person again, and with trade shows coming up in May looks like a busy season ahead, which is pretty cool.

Nishant Lodha 01:09
Really excited about actually seeing people shaking hands. And one day I'll make it all the way to Cincinnati.

Brian Beeler 01:16
California people didn't know where Cincinnati was before the Super Bowl.

Nishant Lodha 01:20
Now we do.

Brian Beeler 01:22
Alright, so we got some good Fibre Channel talk today.

Nishant Lodha 01:27
Let's jump right in. Brian, what's hot in your world? Or in the world of enterprise storage?

Brian Beeler 01:33
I mean, it's, it's sort of a loaded question, right, we've got so many things happening. We've got new form factors for SSDs, which is causing all sorts of excitement and confusion, depending on your perspective, we've got SSDs with all new interfaces, and just as everyone got used to Gen four, well, Gen five and six s, the power SSTs are right behind the corner, which is going to double and double the throughput available. We're seeing all sorts of great things about fabrics as that continues to grow. And I know we'll we'll tear into some of that later. But also, we have the ongoing problems of all sorts of new things that this work from home situation has caused, whether it's accessibility to data accessibility to machines, recovering when we have problems like ransomware, and all these new end users on all these devices have caused all sorts of new vulnerabilities for organizations. So it's as always exciting. But there's so many things happening right now. And even hard drives, I mean, those the guys are up to 20 terabytes now and cranking along, there's so much happening so much opportunity in a bunch of server refreshes coming up this year, which makes the world of connectivity you're really interesting and important too.

Nishant Lodha 02:45
Absolutely Brian I see the word very similarly. I mean, to me, what the pandemic has brought about is definitely a huge dispersion, right. Instead of us all congregating and large shows or within the office, we are dispersed all across working remotely. Not only that, right? Even all our data is dispersed lot of data being created at the edge. And technologies that you talked about right with its next generation PCIe SSDs is high performance, hard disk drives, storage fabrics that we'll jump right into, soon. This all critical to making all of this dispersed infrastructure work.

Brian Beeler 03:24
Well, I mean, the the edge is, is super exciting, but also really complicated, right? Because first of all, it means 1000 different things depending on who you are. It could mean devices, monitoring, industrial equipment. And that data is interesting and important for servicing robots and understanding the telemetry of the date of the machines and the environmental and all sorts of things. But then we've got distributed organizations like banks, and retail and retail, so worried about loss prevention, and how do you protect against fake slip fall injuries and theft and all these things. And that requires an infrastructure. But the point being around the edge is that no matter what your definition of the edge is, the infrastructure in the data is rapidly increasing. And the value of that data is really what it comes down to and is imminently important. And what are we going to do with it? Are we just going to store it and hold it at the edge? Are we going to have accelerators and run analytics at the edge? Are we going to ship it all over the line back to the data center? Are we replicating it? I mean, there's a billion different questions. And a lot of this comes back to regardless of the technology at play. How are we going to move the data and connectivity being such a huge, a huge consideration for these infrastructures.

Nishant Lodha 04:38
Exactly Brian. Irrespective of whether you're running at the edge or within the data center right connectivity, the fabrics that kind of bring everything together are super important. So with that, let's let's dive into into storage fabrics themselves, right that from where I sit or my little port share, I see a lot of choices and fabrics is the new the old or I can say refurbished options in fabrics, and there was definitely RDMA, or more specifically RoCE V2 or NVMe, or RoCE V2 is the most recent standard has been NVMe over TCP, which is just taking NVMe. and shipping it down standard TCP IP, pretty scalable. And also Fibre Channel, Fibre Channel has been around for a while. And several years ago, there was a new standard on Fibre Channel, which makes Fibre Channel capable of transporting NVMe across the wire. With so many options, I wanted to hear what your perspective is because you know, you have evaluated as storagereview.com hundreds, if not 1000s of products 10s of acknowledged transitions, give our listeners the formula or if I can say the algebra of things that they should look at when choosing a fabric.

Brian Beeler 05:52
Well, there's so much there to unpack. And yes, you're right. We've looked at all of these systems, we've looked at all the vendors, we've worked with pure, we've worked with Dell, we've worked with netup a lot. And it's interesting that in the last couple of years, since VMware has formal support for NVMe over Fabrics, built in the hypervisor, that's really been what's changed the game, of course, Linux has had support, and Windows is much more Miss than hit in terms of what you can do there with with NVMe. But VMware makes it real in the enterprise. And I think that's kind of the inflection point for where interest comes in. We talked about NetApp, briefly, but they've offered NVMe over Fabrics in their arrays for years. At this point, we looked at 16 Gig fibre with a 300, a couple of years ago, going over Fibre Channel, and it's really compelling the latency benefits. And I know we don't want to go down all of that, that road right now. But the point being, is that the technology is there, the protocol is somewhat of a question, right? And it really depends in large part of where's your investment today, if you're a big fiber house, for instance, it's very easy to adopt these new technologies. But you've got other vendors, like Dell is kind of standardizing on TCP. So how much flexibility is there in the deployment? Lots and lots of questions. But ultimately, what we're finding, regardless of what we're looking at, from a protocol standpoint, is that there's a big performance advantage. And that's not necessarily a big jump in I ops, what we see it translated to is better latency. And that one's really fundamental to being able to stretch out your storage in investment to make sure that as we're throwing more workloads at these things, is there anything I can do as a current NET app owner, or Dell power store, whatever it is to be able to take advantage of the benefits and latency. And it's normally pretty easy. And we can talk through some of that later, as the latency is the biggest thing, ultimately providing application responsiveness, getting a little bit more outside of the data, or the CPU data path and providing some direct access for better utilization on the application servers. But there's a lot there, it's, it all comes down to in the simplest sense, a big performance advantage.

Nishant Lodha 08:17
For most of my customers, businesses, their data, right. And what I see is that you know, as as lot of these customers have made that journey towards the cloud, right, I see that the business critical applications still continue to be cared for being on prem close to their heart. And when it comes to business critical or mission critical applications, like these databases, financial applications and such they continue to rely on Fibre Channel and the new avatar of Fibre Channel FC NVMe and things that you mentioned right for example, you know, performance and latency or Fibre Channel is like your relevant fully offloaded protocol. Unlike TCP, which and Fibre Channel That's why enables your customers to kind of use their x86 cores to run applications versus spinning them on building packets and payloads. Right when FC NVMe came about. We were very careful to make sure that customers can leverage their existing infrastructure right? There's no rip and replace you know, those those switches from brocade to our Cisco MDs, or the Marvell Qlogic Fibre Channel HBAs. Just make sure they're running the latest software and voila, you got to not just standard Fibre Channel but NVMe Fibre Channel.

Brian Beeler 09:31
That simplicity that I think of what you have and leveraging your current investment is a it may be an undersold point. And by that I mean it when we look at an array and tested is SAN and then test it as NVMe over Fabrics. We're not doing anything different. We're not we're not adding cables. We're not reconfiguring switches, we're not doing anything. All we're doing is re provisioning or provisioning slightly differently in VMware and And we actually did a video recently showing that with a NetApp array of hey, virtualization admins and network admin storage network admins, you don't have to be afraid. And we've done kind of historically as an industry a little bit of a disservice to these guys, right? Because anytime something doesn't work, we always it well, it often isn't networking, but we almost always blame the network admin guy. So if I'm, if I'm that guy, and I've got some uncertainty about what I what should I do this because things kind of work,\I understand it all. Should I do something different to engage on on Fibre Channel over NVMe. And my answer to that is really, yeah. Because there's so much unlimited untapped potential, really, with, especially if you're in a VMware environment, or some other software environment that supports it. in VMware, it's so easy, and you immediately get a latency benefit by enabling this technology. In most cases, it doesn't cost you anything, either via ports or actual costs for licensing. And now your storage array that might have been under duress before because it's more efficient, might be more responsive to other workloads, or be able to take on additional workloads that, that you're looking for a place to go. So that Yeah, I mean, we've got to give the the the network guys a break. And let them be unafraid to use new technologies as it comes available.

Nishant Lodha 11:24
Good point Brian when it when the actual data center kind of gets built or gets upgraded. This is one of the things that come together, right. And the more kind of known and trusted variables you have in this algebra of storage fabrics, the easier everybody's life is right. And the other thing that's often not talked about in general is also the ecosystem around it, right? You mentioned kind of PR and NetApp. But if you look at you know, Fibre Channel, or FC NVMe, and become a pretty strong ecosystem, certainly VMware ESX, and you talked about that, that that's a huge milestone in enabling our customers to adopt FC NVMe or NVMe, or Fibre Channel for their, for their deployments, from storage arrays beyond NetApp. And pure, I believe HPE has a full line of FC NVMe capable arrays. And I'm sure there are others that are missing who have that full ecosystem. And then from boot from a Fibre Channel HBAs I know that when you enable FC NVMe, on the Marvell Qlogic, core fiber HBAs it's not an either are

Brian Beeler 12:30
absolutely. Being able to work in a flexible environment makes a lot of sense. I mean, we do most of our enterprise testing in an ESXi environment, which is pretty typical for the enterprise. But on the Linux side to me, there's there's a lot of flexibility there. And not just if you're picking up the packages and doing it on your own, but there's so many if not, not all, but so many of the software defined packages out there built on top of Linux, and they've offered this support for Fibre Channel over NVMe for some time pre VMware, right. So in the early days of working with these more advanced block storage arrays that would support that it was a little more cumbersome to get operational, because just Linux gives you a lot of flexibility, but also has a lot of extra knobs and buttons that you'd have to understand. But you know, again, I keep coming back to VMware, just from the easy button perspective of saying, Yeah, you can do both concurrently. And again, not have to do a bunch of reconfiguring you're not ripping and replacing anything. And in most cases, you may not even be licensing anything new. So it's, it's it's a net win for for most use cases, I would think,

Nishant Lodha 13:43
Yep. Brian. And today, if I can look at the world around me it, there's a lot of kind of transitions in progress, right. You mentioned on the PCIe side, SSD transition happening that is there's also a lot of transition from kind of rotating media to Solid State, but I still continue to see kind of hybrid arrays have their space. Within the world, I would also expect to see kind of size a small version of the storage arrays show up for to be used at the edge platforms NVMe or fabrics is to me is going to play a very strong role in connecting NVMe to servers and shared storage. If I if my estimate today is just just talking to customers, an analyst is that maybe about five 10% or less of deployments today is NVMe or fabrics but I expect that fast forward couple of years out, I would expect that this will increase to 30-40% in the next several years. So it's it's a huge part of our customers next generation investments and when it comes to mission critical applications. Now their number one choice events Fibre,

Brian Beeler 14:53
It's been historically the one of the if not the most reliable and lowest latency you interconnects available, right. And so that's why it's still around and still event so well understood. And regardless of what team you're on, I mean, almost every environment has at least something for for Fibre Channel and most of its storage related infrastructure related to make sure that that data is passing it at the optimum speed. I mean, it brings up another interesting point in that we did this testing. Most recently, with Dell on an Emmy five, for instance, we use the 32 gig cards, which even if you're not using fabrics, there's still a lot of advantages in the newer technology that's being enabled with these faster interconnects. And when I think about the edge, as we've been talking through this, there's a great argument to be made that we may be able to minimize some of the additional infrastructure at the edge like switches, I mean, you could look at it two ways, we did some testing with no switching with the me five and your 32 gig cards and saw as good of performance as fully switched. So what that tells you is that if you design a small edge infrastructure with a handful of servers, and, and storage that you can get really great performance without having to have a switch over to. On the other hand, if we also make our data transport much more efficient, lower latency and higher speed with these higher speed NICs, again, I'm talking about 32. But I know you guys have 64. And there's all sorts of investments in deep use. And next gen, I'm not even addressing that just looking at 32. When we come from 16. Now we need fewer ports, because we can get the performance out of the array to the servers. And now maybe we don't have to light up so many ports on the switch, if we don't, if we go with a switch versus switch list environment. So there's all sorts of like, you know, we talk about some grand notions like latency at the edge, and it's fantastic, and I love it. But there's also all these other sort of micro implications. And these decisions that end up adding up to quite a big ball of wax as they used to say, you know, if we're talking about lighting up half the ports on the switch, that's big money at the edge times 1000 locations or whatever math you want. And we're not just talking about one SMB here, that's the local accountant in, in eastern Cincinnati, we're talking about scale. And that's where this stuff really snowballs and comes into play.

Nishant Lodha 17:20
People don't talk about it a lot. He's not in the top level kind of conversations. But cost is a huge thing. And like you said, as you multiply cost saving of not using a fibre channel switch on the edge, you can multiply that saving 1000 times like you said, and that's that's huge, you know, there is I know and I've made investments in the Marvel Qlogic, Fibre Channel line to address many of these things are more focused on kind of direct interoperability with our storage arrays, things that you have been doing and have written about. The other big thing also about edge is that these things need to be self driving, right? remote locations, not easily accessible, and it's really hard to manage them. So one of our investments has been around self driving sands, and maybe we should do kind of another deeper dive on that some other day. But the whole concept around this is innovating to make sure that the Fibre Channel HBAs collaborate with the rest of the components in the sand to be aware of conditions in the sand and automatically take actions so that you know, it's it's like a self driving environment that that looks at issues, resolves issues and keeps things running?

Brian Beeler 18:29
Well, yeah, we think of any of the military use cases where they're throwing these on naval ships or submarines. Think about oil and gas exploration. There are so many use cases, I mean, even even space, I mean, we've got a data center in the ISS. Now that's, that's amazing. And as much as you want those astronauts to be able to fix things on the fly. Do we really need them to be experts in storage area networks, or go through an 8000 page PDF on their, on their tablets to service? Why is this port blinking orange, I mean, that's not this, that's not the kind of stuff that that we want those resources doing. And we don't want to have an IT guy in the space station. Now, whether it's that or a branch bank, again, you want the people they're creating value that are offering their goods or services or whatever it is that they do at these locations, not worrying about, oh, no, we've got some sort of packet loss on port 23. On switch four, it's nonsense. So anything that we can do that you guys is Marvell can do as the as the switch providers can do in the software, guys, I mean, let's make this as easy as possible to visualize to prepare without requiring intervention and to also ideally get predictive so that as we see some piece of the infrastructure having maybe about to have trouble that we preemptively get there, and that might be a little harder today, but I think we're definitely had in that direction.

Nishant Lodha 20:00
Definitely Brian, understanding our customers requirements, their needs for scale, reduction of cost making it self driving, that that's what inspires us to innovate every day. So thank you, Brian. And it's been amazing talking to you. Thanks for your perspective around this. And let's get together again and kind of double click into into things that make the sands more self optimizing, and specially for the edge.

Brian Beeler 20:27
Well, we're about to kick off another series of projects where we're using your 32 gig card. So we're going to keep creating the data and next time I'll have you on and we'll, we'll do this again and talk about all the new projects we're working on.

Christopher Banuelos 20:44
Thank you for listening to the Marvell Essential Technology Podcast. As always, please feel free to visit our website to learn more, and we'll see you on the next episode.