Intel

F5 Solutions Running on the Intel® NetSec Accelerator Reference Design

CORPORATE PARTICIPANTS

Lilian Veras
Moderator

Ryan Howard
F5 BD – Principal Solution Engineer

Dirk Blevins
Intel Network and Edge Design – Senior Platform Architect, Principal Engineer

PRESENTATION

Lilian Veras

Welcome everyone to the Intel Network Builders Webinar program. Thank you for taking the time to join us today for a presentation titled F5 Solutions Running on the Intel NetSec Accelerator Reference Design.

Before we get started, I want to point out some of the features of the BrightTALK tool that may improve your experience. There's a Questions tab below your viewer. I encourage our live audience to please ask questions at any time. Our presenters will hold answering them until the end of the presentation.

Below your viewing screen, you will also find an Attachments tab with additional documentation and reference materials, including a number of websites and documents mentioned in this presentation.

Finally, at the end of the presentation, please take the time to provide feedback using the Rating tab. We value your thoughts and we will use the information to improve our future webinars.

Intel Network Builders Webinar Series takes place live twice a month, so check the channel to see what's upcoming and access our growing library of recorded content. In addition to the resources you see here from our partners, we also offer a comprehensive NFV and SDN training program through Intel Network Builders University. You can find the link to this program in the Attachments tab, as well as a link to the Intel Network Builders newsletter.

Intel Network Business partners have been working to accelerate network innovation by optimizing their solutions on Intel technologies. These industry leaders are recognized in our Winners' Circles program and F5 is a Titanium partner. Learn more about our INB Winners' Circle program by clicking on the link in the Attachments tab.

Today, we are pleased to welcome Dirk Blevins from Intel and Ryan Howard from F5. Dirk Blevins is a Senior Platform Architect and Principal Engineer in Intel's Network and Edge Division. He has over 30 years of experience in designing network Silicom platforms and software and has spent the last 21 years with Intel's Networking Division.

Ryan Howard is a Principal Solutions Engineer in F5's Business Development Division. He has over 30 years of experience designing critical infrastructure in the military, aerospace, industrial, FinTech, and healthcare sectors. Ryan has been featured on the Discovery Channel and specializes in energy-efficient large-scale data center and edge deployments.

Welcome, Dirk and Ryan. Thank you again for joining us today. And I will hand over to Dirk to start off. Thank you.

Dirk Blevins

Thank you, Lilian. So, today, we're going to be going through a webinar that is talking about F5 solutions running on a new and exciting reference design product that Intel's brought forward called the Intel NetSec Accelerator Reference Design.

With this said, this is a standard notice and disclaimer about our solutions.

So, what is an Intel NetSec Accelerator? Well, it is essentially a reference design that Intel has developed that allows customers to be able to consume Intel processors in a different way. This is a development of a card that is a PCIe add-in card that has an Intel processor that is plugged into that particular card that allows you to be able to do various functions, such full appliance offload workload acceleration, such as functions like IPSec and cryptography offload, SR-IOV functions. And it allows network acceleration and cloud acceleration functions to be implemented as well.

Some of the key values of this particular solution is that it allows the customer to be able to use code that they know and that they're familiar with. They've programmed with Intel processors for years, so this allows you to be able to take that code, be able to quickly and easily move that code to these particular accelerator cards where it makes sense. There's also key technologies such as Intel QuickAssist Technology that's available for cryptography offload and doing functions such as IPSec.

This processor that we've chosen for this particular solution also has an integrated switch in the product. It would allow you to implement functions such as FastPath offloading so that certain functions and certain traffic doesn't hit the CPU at all.

There's also an integrated ethernet port in the processors that we've chosen in the solution. This kind of capability and this level of integration allows you to be able to actually even fit something – these solutions and Intel solutions in a card that is as small as a PCIe add-in card.

And then finally, this is an independent server. Truly, it's independently configurable. It has an isolated workload boundary to allow you to be able to solve functions like deterministic performance. It has, in addition, independent security domains, so if you need isolation and isolated workload, it's perfect for those kinds of functions.

We launched this particular reference design in partnership with F5 Networks and Silicom. Silicom is an OEM partner that actually you can buy a compatible or compliant card to the reference design. And then F5 is an application developer that are bringing applications and services that will run on top of this particular card.

So, what is this NetSec Accelerator hardware architecture? Well, essentially, we developed a design with scale, performance, and cost in mind. With that said, we developed the design to be able to handle an eight-core and a 16-core version of our P5700 processor family that allows us to be able to have scalability up and down for various solutions. These particular cards – we built two flavors of the reference design, which is a two-by-25-gig card at the lower end, as well as a one-by-100-gig card. We tried to have power considerations when choosing the rest of the features to allow us to be able to place this card into standard existing off-the-shelf servers, where there's power envelope considerations that come into play.

We had the 55 to 90-watt card that's on the left. And then on the right, from those applications that need a little more oomph or a little more horsepower, we have the 16-core P5742 as the base processor that's in this design.

Now, when you look at these solutions or this particular reference design, you can see it has all the makings of a server. At the lower portion of the design, it has an integrated BMC that is on the card with the solution. It has integrated non-volatile memory in the eMMC that's on the middle right side of the solution. Security functions such as the TPM and local connectivity for boot capabilities has its own dedicated memory channels, own integrated ethernet NICs, its own integrated user interface in USB 3 and 2.0. And then it has a PCIe host connector that goes through the E810 NIC that allows you, essentially, to connect to a larger server.

When you combine all this together, it gives you the capability of a fully isolated server that's independently configurable and controllable and maintainable, all within a single small envelope, full-height-half-length PCIe card.

The next question, I think, that most ask is why did we develop this particular accelerator? Why not just use a NIC and more cores, and the other things of that nature? But there's more complexities that get brought about when you look at the solutions and you start considering integration of solutions.

You may have made investments in particular SD-WAN technologies or particular next-gen firewalls. And as the evolution of the function – and the example we're showing here is SASE POP Server – you may want to be able to have those things isolated out from the rest of the SASE functions that may be running the zero-trust, the secure web gateway and those other functions.

So, we really boiled it down to five key scenarios with Workload Isolation, Ease of Integration, Density and Scale, Configurability, and Compatibility as the main reasons why we developed this card.

With workload isolation, I think we touched on it a second ago, but with workload isolation, you have the ability to be able to segment and place applications on this particular card that allow you to isolate things for security reasons. Or if you have something that really needs deterministic performance and gives you the capability to have that deterministic performance directly integrated.

Ease of integration comes down to a couple of different factors. And one of those touches on the compatibility aspect that's down below, where you're just integrating x86 code. So, it's very easy for you to essentially implement this functionality on code that you know and you're familiar with.

There's also another piece to this where certain functions can be standalone appliances running, essentially, virtual versions of these appliances that are now plugged into PCIe slots on those servers, versus being separate standalone appliances. You can choose to service chain those together through ethernet service chaining, or in some of our newer processors, we have PCIe-to-PCIe direct bridging between the different endpoints, so you can be able to use those functions directly in that fashion, making this a very easy element to use and integrate and to make your system easier to integrate.

Density and scale, whereas you may have a standard two-socket server with limitations of core counts going up to 32 to 40, or even 64 cores, as you look into these solutions. And that may be on aper-processor – you might have a two-socket system implementing those functions. This allows you to be able to add more cores, more density in the same footprint, whereas you may have had access to 64 to 80 cores. Now, this kind of a solution can give you, in that same two-rack unit-enabled server, 140-plus cores in these solutions.

And then finally, configurability. As the world moves to more software-defined solutions, software-defined networks, this is again based upon Intel's x86 processing capability, so you have that capability to take advantage of those features.

Now, if we think about this, what's the expected usage of this accelerator? Well, essentially, it comes down to this can be used anywhere a server application can be used. Essentially, we had targeted this for security applications in particular, running on appliances that are shown in the top right and in the bottom left. But the reality is, is a lot of these things are going more virtual, they're going into the cloud. This allows you to be able to take those functions, run those software suites, and do things like what F5 has done, and they're going to showcase in the remainder of this webinar. Ryan's going to take us through a look at some of the elements of what F5 has shown from an application standpoint of NGINX running on this solution.

So, the key thing to take away from this is it's targeted for network applications and network acceleration, but it is a standalone server, and you can use this server to run, essentially, any applications that you want to run.

With that said, Ryan, if you'd like to take it away, please have at it.

Ryan Howard

Thanks, Dirk. So, today, I'm going to go into some of the applications that we can run as example applications on the particular NetSec Accelerator.

One of the common things that Dirk mentioned is scalability. We can take this card and you can put it into existing systems, as we'll talk about in a bit. You have a low TCO. This isn't going to cost you a ton of money to upgrade your systems. It's not like adding another server to the rack. You've got rack space to consider. You've got power to consider. And you've also got durability. When you add more systems, you're having more failure points. You have power supplies that can go out, various hard drives, so on and so forth. And that kind of leads to the durability and the security.

When you have more devices to manage, typically, in different systems, it becomes more of a security challenge, because we have to patch all of those particular devices.

So, as we step through some of the benefits here, one of the neat things about the QuickAssist Technology on-board is it gives us 100 gigabytes per second of IPSec throughput or SSL termination. And what this means is that you can have this card do the heavy lifting of the SSL transactions ahead of, let's say, a web server running in the clear. So, you don't have to worry about running certs on your back end web servers. You just have to worry about running your application.

And when you look at some of the AI performance. This card can increase the AI performance up to 240% over the prior generation.

And as mentioned, they make this in multiple core counts. So, we've got an eight-core count version, a 16-core count version. And this gives you some flexibility and power efficiency and density. So, you don't have to use all the cores at once. You could run a hypervisor that maybe you're running some microservices on a few cores. So, you can scale this out, and when we look at power efficiency – which we'll talk in the next couple of slides – this is not just efficient... even with older systems that are in the market, and newer systems, they will take a certain amount of power, which we'll talk about as to power loss. This will work in legacy systems, modern systems, even edge devices which are quite power-sensitive. So, when you get out to the far edge, you're not going to have much cooling available, and you're trying to minimize your power footprint, especially in modern service provider systems or enterprise banking situations, for example. So, it's very flexible that you can add these cards into multiple COTS servers, which are commercial off-the-shelf.

So, we can do this in a single-slot configuration, or if we're running dual processors, we could add up to six cards, theoretically, in some of the layer model CPUs, because they have more lanes. So, this gives you a core count of upwards of 176 cores in a 2U system, for example. So, you've really increased your power density, and you haven't increased the power draw on that particular server, nor have you added extra rack space. Which, in today's environment, we're trying to keep the systems as dense as possible, because every U of rack space costs money, especially if you're leasing it out at a large data center.

So, I alluded a little bit to power efficiency and density. When you run a particular server, just when you look at the power efficiencies of the power suppliers, you'll sometimes see platinum and gold. And really, what that's talking about is how efficient that is converting from AC to DC, which is what the system is trying to offer. And typically, you're going to see an 80-watt efficiency – sorry, 80% efficient power supply. Which means on a thousand-watt system, you're going to be losing 200 watts right off the bat to heat. So, if you could basically take – instead of adding four more servers, add three cards to your server, three IAONIC cards, you would be eliminating 600 watts of just wasted power, plus the power those systems take to run their primary processors.

And as the use case – that we'll walk through here shortly – you don't often need sometimes a full server. You can put these cards in and run a lot of your dev and staging loads on it to make sure something doesn't break when it hits production.

So, in summary here on this slide, you really do get the power efficiency and density without the cost, when you're spending over 10 grand, 10, $20,000 buying additional servers. Why would you spend $60,000 when you could buy these particular NICs and use the exact same server that you have now? Especially with supply chains looming, orders taking longer, this will give you an instantaneous boost in your systems.

So, as we spoke, on the lower end, this takes 55 to 71 watts. Some of the higher-end SKUs can take a bit more. But under 75 watts, you don't need a GPU power connector off your board.

It looks like we lost the slide. There we go.

All right. So, that's an advantage because a lot of – you'll hear the term "SmartNICs' thrown around out there. And some of the SmartNICs are very power-intensive, even though they may be running ARM cores. This particular card is x86, so typical applications will run on this. They don't have to be compiled for ARM, which is typically problematic. And you don't need to run a bunch of obscure GPU connectors to provide auxiliary power beyond the 75 watts the PCIe slot can put out.

So, the NetSec Reference Accelerator Design enables the ability to have, essentially, blade servers. This is – so you can think of them that way as adding additional blades, but they're in the form of a PCIe card. And every particular server – this references a 1U, but it could be a 2U server, it could be a 1U server depending on how many slots you have left. But theoretically, you could throw up every single remaining PCIe slot that you have and get 16 cores plus 100-gig NIC plus QAT in that particular slot. So, there's a lot of advantages there.

So, we're going to talk about the use case here of running NGINX App Protect. It's our web application firewall. And we could run this natively on the IAONIC card.

So, one of the advantages of NGINX is we can use QAT acceleration. As we spoke to a little bit, what is QAT? QAT is Intel QuickAssist Technology. We can take the QuickAssist Technology and avoid the heavy lifting that a CPU normally does, basically translating the crypto side, or even on a VPN, and we can send that off to this QuickAssist engine. And it will process this for us as a sideband to the CPU. So, we're really lowering – significantly lowering, and we'll look at the numbers here shortly – the overhead on the CPU, which raises the number of transactions per second we can do on the SSL side.

So, when we look at the general numbers in an average system, we're going to see, roughly, a 10x improvement in the number of transactions per second we can do with QAT. By default, this is so horsepower-intensive, you might get five to 10,000 transactions per second just using the raw CPU horsepower, but you're maxed out at 100%. And it doesn't leave much overhead for analysis or anything else.

With only assigning even 12 cores from a 16-core processor, we're up around 92,000 transactions per second on a single card. So, you can do the math and add this to – add several cards, let's say, three cards plus your primary system, it starts to turn it into a system with quite a bit of horsepower.

And so, we start talking about what applications – what can we do with this? And a lot of times, you'll have your NGINX as your software load balancer, your API gateway, your NGINX controller. You might have some monolithic apps, some API endpoints, some microservices or, let's say, Kubernetes pods. It could also be other VMs, it doesn't have to be Kubernetes. And one I've mentioned in the past quite a bit is – and still to this day – is you have infrastructure folks, DevOp folks, NetOps, security, when an app is rolled out, they start and it breaks. And some developer will say, "Hey, I want to add this new function", but they haven't rendered it through any sort of testing with the outer web application firewalls or even basic firewalls. And so, they might write a particular code that has various violations that trip the web application firewall.

And so, it's nice to be able to separate that out into a dev environment, which is an isolated sandbox, not real data, staging which typically has, let's say, three/four-day-old data if we're talking healthcare or banking, and then a production environment. And the production environment could be the main actual system running the workload.

So, what's nice about this is you have three separate 100-gig NICs that you can use in the dev, staging, and prod environment. So, that gives you plenty of bandwidth. You're not wasting a PCIe slot by just adding a QAT card. You're essentially adding a 100-gig NIC interface. You're adding 16 cores. And you're adding QuickAssist to every card you have. So, you're giving actually a fairly decent environment to the dev environment, and they can run through the exact same profiles that we would run on the mainstream production.

So, basically, they can have their sandbox and see if it breaks. And this is nice for security because security can come in and make sure everything is triple-set as far as the performance and the security policies, and the app is simply working.

And then as you move into staging, you can test this out with real data, make sure it functionally works.

And then finally, you can make the switch over to prod with a lot more confidence that the code is going to work and that you don't have to roll it back and have interruptions in service. Because the worst thing that can happen when you're live is to go down even for – minutes of downtime could cost millions of dollars. Or in the case of healthcare, it could cost lives. Certainly, it would affect scheduling, and during critical times, let's say, even in the retail sector, if you're down for, let's say, even 10 minutes, you could be talking orders of magnitude, especially on holidays and events where you may be trying to process a massive amount of sales transactions.

And again, that's another application beyond dev, staging, and prod is you can scale out to these cards to add more horsepower in those events. Especially what we're seeing in the banking environment, the FinTech is that we're seeing large flexes in the amount of capacity planning they need to have for customers.

So, it will be – let's say, a news event kicks off. It used to be 30%, 40% overhead they need to have available. Now, we're talking two to 400% extra capacity you need to have available. And really, as we talked about earlier, you don't want to do that in the form of adding four more, 10 more servers at $20,000 a piece, it starts to get expensive from the heating, cooling, rack space standpoint, and management as well as security.

So, with this, it gives you a much simpler solution to scale out as well.

So, one of the other use cases that we mentioned is also edge. You could put this out in a lightweight edge box. A lot of times, they may have a PCI slot, but they may not have a lot of horsepower. Again, you could have this doing firewall functionality, WAF functionality out at your edge, let's say, a branch office, hospital, clinic. This is a perfect, low-power solution for that.

So, as we step through this – this is kind of a visualization of an isolated sandbox where, on the switch, we can put this out in its own private Idaho. It can be on its own separate VLAN. It can be routed differently. You could say, "Hey, it doesn't have access to the internet at all". And really, you can start to sniff that traffic to see if there's a problem with it. Then you have your staging, again, it's kind of an interim firewall where you can say, "Well, this can partially get out to the internet". Maybe you can pull traffic down or do API calls back to our database, but it can't get out to the internet or public world. And then you have your production. You could run prod on these cards or you could run the production on the primary CPU.

So, when we start talking – what can we do with this? Well, we can – we want to let the good traffic in. You can saturate a common off-the-shelf server – or commercial off-the-shelf server very quickly with common attacks. You could absorb time slots, for example, in any sort of scheduling application. You could say, "I want to fill up all these deliveries with fake orders". We can start to detect that with our engine and say, "No, that's a bot versus a real person". And we can deflect bad actors away, especially if they're trying to do this for cross-site scripting injection attacks, fuzzing, zero-day, a path traversal, man-in-the-middle type of attacks. We can say, "No, we want just the good traffic to come through". And keep in mind, it doesn't take a whole lot of traffic nowadays with these high bandwidth 5G phones, it might take an order of 10, 20, 30 devices to take down a website.

So, one of the things you can also do with these particular IAONIC cards is you can use them as your DMZ. So, you can say, "Well, I want to send all the traffic to these particular cards first and see if the traffic – weight the reliability and then accelerate them to a different path to show... to send out, let's say, the main compute".

So, when we talk DevOps integration, NGINX is obviously wonderful. It runs on the majority of the world's web servers. We have integration with Red Hat OpenShift, Kubernetes. We can graphically display with Grafana, Datadog, Nagios. And we have all sorts of existing open source tie-in. You could do this and use these cards at very, very low cost to your infrastructure, which everybody cares about right now.

So, back to the TCO story, instead of buying a bunch of servers, these cards will integrate just like your regular x86. They're not on an ARM processor. They look like your regular server to your hypervisor. And you can just run your normal operations as you normally would.

So, we have a very lightweight WAF. So, our lightweight WAF, as I mentioned earlier, will stop command execution vulnerabilities, cross-site scripting, information leakage vulnerabilities, SQL injections, authentication attacks. Basically, you could exhaust someone's credentials – let's say, there was only three times that they needed before it was an absolute reset, and now you have customers that can't even log in to the website, because someone's reset that. So, we can stop those sorts of attacks. And also, things like buffer overflow vulnerabilities.

And then your servers in the back end, which could be an app server, web server, app delivery controller, API gateway, container, a whole variety of microservices which could be running on the cores on these particular cards, or on your main host, or elsewhere in your network. We can protect those with our NGINX App Protect WAF.

So, this is all Layer 7 DoS security. So, we can protect things like Slowloris, slow read, slow POST, flood Attacks. We can do TLS fingerprinting, so on and so forth. So, it's pretty comprehensive. You can check out the NGINX App Protect website, which will give you more information as to what we can protect against and how we can integrate into the solution.

And lastly, if you have any questions, please contact your F5 or Intel representative and they'd be happy to go over this product with you.

Back to you, Lilian.

Lilian Veras

Thank you, Ryan and Dirk, for sharing such great information with us. We do have a few questions that have come in while you were presenting, so let's get started on those.

The first question. "Does Intel plan on developing more reference designs like the Intel NetSec design?" Dirk, you're on mute.

Dirk Blevins

Let me come off mute.

Lilian Veras

Yes. No problem.

Dirk Blevins

So, the reality is is that we don't intend on this being a single isolated design. We absolutely do plan on expanding the product portfolio. We do not have details to share at this point in time on what that expansion will be. But at this point in time, we do plan on expanding the portfolio.

Lilian Veras

Awesome, thank you. Question number two. "Is there any more information on which OEM or ODMs will bring the product offering to market?"

Dirk Blevins

Okay, that's coming at me again. The reality is is right now we don't have any more detailed information on that. What I think I can share is that there will be several – there are definitely several others that are in process of developing products to come forward. And I hope within the next quarter or so we'll be able to share more details on who those particular partners are and how those products can be utilized.

Lilian Veras

All right, thank you. A question for Ryan. "Why would using these cards save power compared to adding more servers?"

Ryan Howard

Sure, no problem. So, typically, as I spoke about a bit in the chat, each server has inherent inefficiency. There's wasted power based upon the power supply. So, if your power supply is 80% efficient and it's a thousand watts, you're losing 200 watts just simply out to heat. And we start to scale that out to many, many servers, that why you need to cool data centers, because not only are BTUs being put out from the loss just right away from the power supplies, but then you have all of the heat generated from the CPUs and the internal components.

So, these cards are very lightweight, they produce very little heat. So, when you add these cards to the server, it's just a mere 75 to 115-watt power increase, which is typically handled by your existing power supplies. You've already lost that 200 watts. So, you're not going to lose any more. You're just adding an additional clean load to the power supply.

And you're also not adding – which a lot of people forget about – is a BTU load to your HVAC systems. So, you're always worried about cooling in any sort of data center. And that is where the bulk of the cost is, is running large HVAC systems to essentially air cool these servers, and that whole process is not very efficient.

So, again, going back to – you're not taking any more power, and these cards are ultra-efficient, so they're really not adding to your BTU load much at all. So, you're just adding a clean, let's say, 74 watts to the system, multiple cards, 150, so on, and so forth, which is usually well within your power supply range. And then compared to the BTU load that you would add – let's say, you needed to add four more servers versus adding three, four cards, those four more servers, let's say, you're adding 800 watts plus the CPU consumption. So, you're probably up around 2,000 watts of power, two kilowatts at least, maybe upwards of four kilowatts that you're consuming. Plus, you're adding a CPU load of four additional servers.

So, it starts to become very expensive very quickly, as well as not to mention the rack space cost on a per-U basis.

Lilian Veras

Awesome, thank you for that detailed explanation. Another question here. A member of the audience is asking if it matters which hypervisor or OS he uses with the IAONIC?

Ryan Howard

Sure, I'll take that one. No, you could use pretty much any hypervisor, any OS. You could use KVM, VMware, OpenStack, OpenShift, Robin.io, any of the new flavors out there. It's an x86, so that's the key point here. A lot of the SmartNICs and DPUs out there that are advertising, you can't run much on them. They don't have a ton of horsepower in their ARM. So, you have to recompile your applications, your applications may not even run on the ARM. And even if they did, they wouldn't have a high enough clock frequency to actually run the application.

So, this is putting a server-class card and a server-class CPU into your system.

Lilian Veras

Awesome. Another question here. "Can the CPU on the IAONIC run multiple VMs? And if so, how do I assign a 100-gig NIC to each VM?"

Ryan Howard

Sure. So, yes, you can actually run multiple VMs. You could carve it up to run a VM or microservice on each particular CPU. These are actual cores that we're talking about, so they're not hyperthreading. They have their own memory available to them, which is ideal. They have their own storage available.

And so, how do you connect the NIC? The NIC supports what's called SR-IOV. And SR-IOV allows you to clone the cards so that your VMs think they have their own dedicated NIC. That's all – it's called a Virtual Function. So, you can assign that virtual function, you could create 500 of them if you wanted to, and assign those to each particular VM, if that makes sense.

Lilian Veras

Okay, awesome. One last question – one more just came in. So, one question for Dirk here. "This product looks a lot like SmartNIC that others have been talking about, what makes this different and would you call it a SmartNIC?"

Dirk Blevins

Oh, boy, that's a bit of a – a little bit of a trick question. But Ryan kind of alluded to the answer to this just a second ago. And I'm going to go back and lean on what he just said that, in reality, it does have a lot of elements that look and feel like a SmartNIC. There is acceleration offload capability, there are processors there that allow you to program the capability of the card. There's a lot of things that do give it the look and feel of a SmartNIC. So, I guess, technically speaking, you would call it a SmartNIC.

Now, with that said, we really had a very targeted focus for doing this. Network and security applications and kind of what we coin as the "bump in the wire" application. We definitely have a lot of – I think we have a lot of advantages and the fact that there's a lot of code that exists that will just run on these solutions. And we're seeing uptake from the customers that are looking at the product to be able to do it.

I'd say Ryan would tell you that porting their applications that they've started to port have been fairly straightforward to do on these solutions. And I don't think you're going to get that with what people have coined as "SmartNICs" in the marketplace today.

Ryan Howard

And I'll add to that. What's very neat about this solution is it's a little bit different than a SmartNIC because it's a standalone environment. So, you could PXE boot if you wanted to, and create this as an OpenStack endpoint. You could make it a standalone system. So, not only can you see the card from your host system, which is typical with SmartNICs, but it can act as a standalone compute node if you want to because it has storage, it has long-term storage on it, and there's no reason that you couldn't just PXE boot off the 100-gig NIC and have a standalone system, completely isolated from the rest of your system.

And we alluded to that in the dev example – dev versus staging versus prod environments. This dev environment could be routed even on the switching side to its own separate world.

Lilian Veras

Great, thanks to you both. We do have one last question. A member from the audience is asking, "Will DPDK run native as well?"

Dirk Blevins

Do you want me to take that, Ryan, or do you want to?

All right, yes, so – yes, DPDK will run native on the solution. In fact, we've had folks that have looked at trying to – well, not looked at, they have taken and ran like VPP with StrongSwan running natively. They've taken DPDK and ran some of the example apps that would run on our standard P5700 family processors.

So, the software just runs straight out of the box as if you're running the software on a virtual appliance or anything else.

Lilian Veras

Awesome. Well, thanks to you both for such a great presentation. This concludes our webcast.

I will ask our members who are attending the live webcast to please do not forget to give our team a rating for the live recording so that we may continuously improve the quality of our webinars.

Thank you, Ryan. Thank you, Dirk. And this is all for today.

Dirk Blevins

Thank you.