

AdvancedTCA and beyond in '007: Licensed for killer apps?

Q&A with John Fryer and Shlomo Pri-Tal

As 2006 comes to a close, it seems appropriate to consider what 2007 and beyond will bring for the AdvancedTCA ecosystem.

OpenSystems Publishing Editorial Director Joe Pavlat spoke with John Fryer and Shlomo Pri-Tal of Motorola.

2007

OSP: John, you recently addressed the AdvancedTCA Summit and had a chance to take a look at the state of the technology. What did you find particularly interesting?

John: Increased interest around the software aspect and platforms in general. There's interest around the Service Availability Forum (SA Forum) and movement from interest to adoption in the marketplace.

OSP: Who do you see as the main suppliers of SA Forum-compliant middleware?

John: Aside from us, the two major players out there are GoAhead and OpenClovis.



I think the players recognize that it is a big step to go from a software solution for high availability to a hardened platform that encompasses both hardware and software. Clearly that path is a big part of our strategy with the Avantellis platform. The industry wants to adopt a fast and inclusive solution across a very broad range of platforms, not just across COTS platforms, but across proprietary platforms as well.

OSP: Management on the equipment provider side seems to understand that they need to focus on intellectual property, yet we have 100 years of engineers proclaiming: "We have to do everything ourselves to ensure quality and reliability."

Are cultural problems going to keep the rate of AdvancedTCA adoption artificially low for a number of years?

John: There are a couple of perspectives on that: If you run back the clock say three or four years ago when AdvancedTCA was just starting to appear in the marketplace, it had a promise, but really had no market, either existing or evolved. There were experiences from CompactPCI and PICMG 2.16; certainly the 2.16 platform we could assess as being moderately successful at best.

Now if we fast-forward to today, you have commitments to AdvancedTCA technology from virtually all the network equipment providers, the exceptions being Ericsson on the wireless side and Cisco on the wireline side. They are market leaders, and it is natural that you will have reticence among the market leaders to adopt a new technology, because if you lead the market, you think you are doing it right.

We are definitely past the point of: "Should we adopt AdvancedTCA technology?" and have moved onto: "How do we adopt it and when?" On the carrier side we are seeing RFPs with a requirement to implement AdvancedTCA capability.

OSP: Why do the carriers care?

John: If you look at a lot of the network architectures out there today, every time they want to implement new functionality, it essentially means they need to add a new device. So it makes every aspect of the business more expensive: You have to consider the complexity of network design, installation, and maintenance.

It is not just the capital expense, it is the operational expense – they have to keep spares for N different devices. Now it is N+1 different devices.

OSP: *So carriers got the religion that AdvancedTCA would reduce their CAPEX and OPEX?*

John: Yes. That's the promise. If you look inside the equipment manufacturers, legions of engineers continue to believe they can do the job better than AdvancedTCA can do it, and in all probability, if they set out to design a specific piece of equipment, they can do it better ...

OSP: *... and have it ready for market in three to five years.*

John: That is not the name of the game any more, the name of the game is all about time to market and rapidity of these services and being able to collapse multiple functions together. One of Shlomo's favorite sayings three or four years ago was: "Networks into boxes, boxes into blades."

What the industry is all about today is collapsing functionality and combining functionality in new and different ways, some of which cannot be foreseen today. It's important to have platforms that are good enough, and flexible enough, to do the job rather than having the perfect platform for a specific network function.

Shlomo: I think the adoption of AdvancedTCA is a done deal, there is no question. I think the challenge we have now is to deliver. And when we say, "deliver" we don't mean a blade, we mean platforms – that is what customers want from us. Carriers' time to market is to launch services; they don't give a darn about the hardware. They see the industry widely adopting AdvancedTCA with all the other things necessary to make a platform as a means for them to be able to launch services quickly.

OSP: *How long will AdvancedTCA have sufficient force and power to be a major platform in the market?*

Shlomo: It all depends on two things: One, what can you accommodate from a fabric perspective, and two, what can you do from a heat perspective?

If we as an industry can develop very high-speed fabrics and accommodate 500 W and above per blade, it could survive for a very long time.

OSP: *We are going to have to support 10 gigabit per pair technology?*

John and Schlomo: 10 gig [per slot] is not the issue. The issue is to go beyond 10 gig, to have 10 gig times four per slot.

OSP: *With 10 gig per pair, as opposed to 3.125 gig per pair, which is where we are now, heat is going to be a problem because of the need for higher-powered processors to handle the traffic. Best practices in VME and CompactPCI were around a watt per cubic inch power dissipation in the subrack. AdvancedTCA made it to about 1.25 W per cubic*

inch. I don't think we're going to double that amount easily with just forced air.

Shlomo: It is reasonable to say that the network transport core is going to be very specialized; don't worry about it, and don't try and capture every last application, just stay where the volumes are. You can survive with 10 gigs per pair [up to 40 gig per slot]; you do not need more than that, not for 10 to 20 years.

But the core is a different issue – they are looking at multiple 10 gig and 100 gig pairs – forget that, just don't go there. There are always going to be fringe applications where it is highly doubtful a standard environment is going to be adopted.

The other thing to remember is success depends on the ability of the industry to deliver reliable products. That is why SCOPE [Alliance] is so critical, because they tell you what they want, which we did not have before, and [the Communications Platforms Trade Association] CP-TA is so important because it assures that we can deliver what they want – and that it works. (See *OSP Editor's Notes* February 2006 *CompactPCI and AdvancedTCA Systems* compactpci-systems.com/columns/Editors_Notes/pdfs/2006,02.pdf. SCOPE Alliance is an association of the major network equipment providers.)

If we do not succeed with these two activities, and if we do not make them the centerpiece of the ecosystem to deliver product to the NEPs, AdvancedTCA is not going to deliver on its promise.

John: So building platforms to the SCOPE requirements is also the way we will bring down the remaining barriers with the equipment manufacturers.

OSP: *What percentage of the overall market will the profiles that SCOPE is developing be able to address? 10 percent? 50 percent?*

John: That is a challenging one to answer.

Shlomo: It depends on how you count it. I don't think we will ever get into the core, in spite of the fact that maybe we will have equipment that can do that. It is an area that the NEPs are not going to relinquish for a long time.

Now as to access, AdvancedTCA is not going to go into access – it is too big, too cumbersome, too expensive. MicroTCA may very well, unless we overcomplicate it and burden it with features that may not be necessary. So if you say that 50 percent of those two at the edge and the access is open to us, that is a pretty big number.

John: And I think AdvancedTCA is a done deal for anything pertaining to a telecom-related control plane, particularly where you need a blend of computer-centric applications with packet processing or media processing capability.

OSP: *Speaking of packet processing, with the demise of Intel's network processor group ...*

Shlomo: Network processors are a thing of the past. It is a different programming model, nobody understands it, and nobody can

program it efficiently. What our customers are telling us is that they want to do both control and data plane using the same architecture. So what you are going to see is the increasing importance of multicore processors. MIPS64 in this case will be used – you will see some limited success from that, and eventually you will see high multicore processors from Intel. Whether or not they are going to be successful in this particular space remains to be seen, but they are not going to give it up.

OSP: *What about FPGAs, do they play a role?*

Shlomo: For specialized functions, probably, but when it comes to just mainstream packet processing I do not see why anybody would use anything but multicore general purpose processors particularly when the various engines are attached to the multicore processor itself.

John: Some things like FPGAs may well play a role where an equipment manufacturer decides to build a blade. Equipment manufacturers will continue to build blades for areas that are very specialized, where they feel they have a particular expertise they don't want to give up – and that is a model that sits along very well with an off-the-shelf communications server type platform.

OSP: *How is convergence of IP and legacy TDM influencing the adoption of AdvancedTCA?*

John: We are seeing markedly increased interest around pure IP applications, that is, packet-based applications. The need, and the legacy requirements, are still there, but diminishing rapidly. They are more down to the interface level rather than the transport level. For example, a year or two ago there was more talk of: "How do we switch TDM streams through AdvancedTCA devices?" There were all sorts of ideas about how to do that, and quite a bit of work was being done in that area, but I think we are seeing a diminishing interest in that. Now pure IP-based applications are starting to come to the fore. Just look at the growth of VoIP over the last year, and the writing is clearly on the wall. So there is less of a requirement to support TDM-type stuff and more of a focus around IP.

One of the questions popping up with regard to Shlomo's comments about going beyond 10 gig is: "How can we tackle the IP transport plane, or the Ethernet transport plane, and what should do about that?" We can't do that in an effective way with AdvancedTCA; we can do it in a limited way, but that would be one of the next interesting places to go."

As Shlomo said earlier, the challenge is how to evolve these platforms over successive generations of technology until there is a general consensus that you cannot do any more.

Multicore technology has come along at a very opportune time because it changes the dynamic. Much of the path forward seems to be bigger, hotter, more powerful single processor Xeon-type technology, but we have been getting a lot of push back saying: That is not the way we want to go, we want a profile that fits within the 200 W envelope. Silicon vendors are being pushed to come up with that. Now solutions are appearing, and Intel is on that bandwagon as well.

OSP: *I suspect that absent new developments in device physics, which we have not seen in 20 or more years, just the heat, as well as the sheer number of transistors and the amount of memory required, are going to create obstacles at some point.*

Shlomo: Yes, but at the same time the point that John is making is very true, we have to push back on suppliers.

OSP: *Speaking of fabrics, is everything other than Ethernet essentially dead?*

Shlomo: Yes, for AdvancedTCA. For AdvancedMCs and MicroTCA, I do not consider PCI Express to be a fabric, but it probably would have a use. Also RapidIO, because of its use in DSPs. The mainstream market will be Ethernet. What we should do as an industry is make Ethernet have the capabilities to take us into next-generation systems, rather than inventing new fabrics.

OSP: *Which tells me we have to get to 10 gig per pair in a couple of years.*

Shlomo: We have to do whatever we need to do to assure we have the runway to stay with Ethernet for a long, long time. You are not going to have any other fabric that delivers the commercial benefits Ethernet does. It's going to be adopted by the enterprise market in general. It is a done deal; let's go make it perfect.

OSP: *One of the elephants in the closet is Huawei. Their business model has always been to try to develop some of their products around standards, so they can fuel their export market. Granted, they do not have much presence in North America and may never, but they are pretty big in other parts of the world. Do you see Huawei weighing in?*

John: Well, I think they are already there. I mean we know they have significant investments in AdvancedTCA. Huawei's presence in general has driven a lot of manufacturers toward AdvancedTCA whether they like it or not to get to the price point that enables them to compete with Huawei. They are a factor on everybody's radar, including our own as a potential customer but also as a potential threat in the marketplace.

OSP: *And looking ahead to 2007 ...*

Shlomo: 2007 will be the year of assuring interoperability of products. While the overall trend is for integrated platforms, there are still enough customers who want to buy building blocks and integrate them, and we have to assure they have an absolutely superb experience.

OSP: *The equipment is moving out into carrier deployment?*

Shlomo: Absolutely. We see that today.

John: So volume building is one of the things we will start to see in 2007, and will even start to see in 2006. The challenge is delivering full platform to 5-nines. I think we have gone through some of those growing pains this year, and they should start to pay dividends in 2007.

OSP: *Were there more growing pains than you had anticipated in 2006?*

John: The whole industry discussion around thermals, and what it means for a platform to be 5-nines, and the integration of the software with the hardware is a substantial effort.

We will start to see the benefits in 2007. It is a lot easier to walk into somebody's office and say: "Do you want to take these building blocks and use them yourself, and you can layer on some software, and it will all work." It is a major step to go from that to a hardened 5-nines platform. I think we will see much more of a trend toward that kind of decision in 2007 as these platforms appear in the marketplace.

OSP: *What will the issues be for AdvancedTCA five years from now?*


Shlomo: The ability to continue evolving in a backwards-compatible manner.

John: Once we get past the basic interoperability stuff, that is what the industry has to focus on.

Shlomo: Keep the connectors, keep manageability. Keep backwards compatibility on the server to the extent that you understand that at some point there might be a departure, but it would be very far in the future.

So, yes, backwards compatibility is crucial. That is what has kept VME going for 25 years now. It is not the technical features or anything else, it is the fact that you do not reset the market every time you come up with something new – you leverage what you have already built, you leverage the volumes.

OSP: *What type of a lifetime do you see for AdvancedTCA in terms of being a mainstream platform for telecom?*

Shlomo: I see no reason why it should not last as long as VME. 



John Fryer is director of technology marketing for Embedded Communications Computing within Motorola's Networks business. He is responsible for determining market trends and future customer requirements to develop strategic directions of Motorola's AdvancedTCA and MicroTCA based communications servers.



Shlomo Pri-Tal is president and chairman of the Board of Directors of the Communications Platforms Trade Association (CP-TA). He is also the director, Group Technology Office, for Motorola's Embedded Communications Computing business. He oversees Motorola's participation in standards development organizations and industry consortia.

OSP: *Are you optimistic about MicroTCA?*

Shlomo: If MicroTCA is an environment that allows us to go to the access layer of the network in a very efficient way, then it is going to be successful. If on the other hand we start complicating it by making it dual star redundant, and 5- and 6-nines reliable and essentially dilute MicroTCA in relation to AdvancedTCA while burdening MicroTCA with capabilities that are not necessary, I am less optimistic.

OSP: *I happen to agree with you on that. I worry about a real chicken and egg problem with MicroTCA. The costs of an empty chassis, with its specialized power supply, individual control lines to each slot, and a whole MCH are not inherently a high expense if someone is willing to invest and bring out silicon. Right now it is fairly expensive. At the same time, the silicon manufacturers are going to have to see MicroTCA volume as attractive before they are going to invest in it with chips. My worry is that MicroTCA is stuck at the moment in this area where it is expensive and vendors are concentrating on high end, fully redundant, duplex MicroTCA. And at the high end, it competes too much with AdvancedTCA, especially as AdvancedTCA volumes increase.*

Shlomo: What is interesting is that we see quite a bit of demand and interest in MicroTCA from non-telco applications, particularly defense and aerospace and medical. We see substantial volume coming up on MicroTCA; there is a lot of excitement in the marketplace about it.

John: I think one of the issues with MicroTCA in talking with some of the telecom guys, is that if you move more and more to the edge where MicroTCA would naturally fit, those become more and more cost-sensitive applications, so I have had people say to me: "Look, we spend every minute of every day taking every cent we can out of the cost of our products, so what you have to be able to show us is that you can deliver a general purpose platform equivalent at any cost point to what we can do today." That is a significant challenge given where we are today in the evolution of MicroTCA, and it speaks to Shlomo's point of: Be careful what you burden MicroTCA with in terms of the telecom-type applications.

The telecom market is still developing for MicroTCA; it is potentially so flexible, and there are potentially so many form factors being talked about. There are myriad packaging options driving to different markets, each of which has a different price point, and it can become a complicated equation. However, MicroTCA volumes will be driven by industrial, defense, aerospace, and medical applications so the telecom folks may be able to take advantage of those economies of scale.

OSP: *I get widely varying thoughts from people on how much of the AdvancedMCs market is going to be for MicroTCA. Do you see AdvancedMCs as representing 2 percent of AdvancedTCA volume? 10 percent? 20 percent?*

Shlomo: I don't know the answer, and the reason that I do not know the answer is that it seems doubtful that the same AdvancedMC can accommodate both environments [AdvancedTCA and MicroTCA], especially if you design it for dual star and things of that sort. I believe that MicroTCA will drive the volumes of AdvancedMC.

OSP: *Which means the AdvancedMC architecture might change to be more cost-effective for MicroTCA?*

Shlomo: Which makes it less cost-effective for AdvancedTCA. Now you reach a tradeoff between laying capabilities in the blade versus having it in a plug-in module.

For more information, contact John and Shlomo at:

Motorola Inc., Embedded Communications Computing

2900 S. Diablo Way

Tempe, AZ 85282

602-437-6701

john.fryer@motorola.com

shlomo.pri-tal@motorola.com

www.motorola.com/computing

For Single Print Only