An Argument for Standard Architecture Development

"Will companies use proprietary or standard's based platforms?" The telecom field is not moving today from traditional media layer platforms to switches, meshed switches and meshes, this is only something happening currently in PICMG and VITA. Most network routers and telecom switches have been switched fabrics for many years. They are just proprietary designs.

On December 30, 2002 the Executive members of PICMG adopted the Advanced TCA specification. The Advanced TCA specification is 430 pages long and contains extensive electrical guidelines for a full mesh backplane. This 400+ page document is a comprehensive technical primer that should be required reading for every system architect or engineer approaching the design of a high performance telecom switch.

If it is correct that networking and telecom equipment manufacturers have little incentive to migrate from custom platforms to standard architectures, why else might a standard architecture interconnect be of interest to design engineering firms?

There is no doubt that there will be many projects that will involve implementing a full square mesh or a dual dual star with multi-gigabit differential signaling. Accomplishing it within the requirements of PICMG 3.0 would be an exercise that would be re-usable on subsequent customer supported design tasks and serve as an open interconnect platform model for customer design seminars, and technology discussions. A public effort can serve as a focus around which to discuss common technical issues with a diverse group. After completing a proprietary design effort you have nothing that can be shared with the next customer.

There is a limit to the practical size of a full square mesh fabric because if there are sixteen slots then each card has to have n x 16 connector segments where n=the number of segments required per link. The routing density and number of layers to accomplish the interconnections is huge. PICMG 3.0 will require around 28 - 40 layers for a sixteen slot full mesh with 8 pairs for each link and supporting 3.1 Gb/s bit streams for each pair.

The magnitude of the task is not trivial and the sooner you have done it the first time the sooner you will be able to accurately bid on doing it for a customer. It is important that the design team that tackles this task be the same team that will be available to work on customer implementations. A number of design teams have already accomplished a similar task on a smaller scale with 2 row ERmetZD connectors. Quite a few have done it as well for 2mm HM connectors and many have done it for HSD connectors. Today, however, the ERmetZD connector is the right connector and the layout requirements will make the ERmetXT attractive for a following generation.

Another important question is to ask yourself is "What comes next?" The question of midplane cross connect architectures poses difficult technical questions for a differential connector, however, a midplane cross connect architecture may offer some important advantages. There are signal integrity advantages - allowing more space for the routing by using the "Z axis" space behind the backplane to space out traces for better performance. There is also the ability to use any given connector on a slot card to connect with any other slot. This could allow the same non-blocking connectivity that the full mesh architecture provides with only half the connectors on each card. At present, this is an academic because the connectors to support this architecture do not exist. At least they do not exist to support a three-gigabit differential signaling. Since we are talking about the future if such a connector was planned it would have to target 10 or 12 Gigabits/second.

In any current discussion of differential, serial point-to-point meshed architectures the advent of multiple silicon I/O semiconductors incorporating active, dynamic signal conditioning should certainly be mentioned. We should remember that Hayes and other modem manufacturers adopted increasingly sophisticated signaling techniques to move from 300 and 1200 baud to 56K baud data rates, over the existing unshielded twisted pair structured telephone network. Today Ethernet technology is preparing to move rapidly from 100 megabit/sec to 10 gigabit/sec also over similar twisted pair wiring. Also today, almost every ASIC supplier offers some collection of advanced, integrated suite of signaling technologies which include various dynamic signal conditioning techniques such as back channel feedback control, dynamic equalization, active de-emphasis, and programmable multi level signaling.

An Argument for Standard Architecture Development

These approaches can be expected to super charge traditional serial backplane architectures when the communication industry has recovered sufficiently to support new development activities.

Also, what about other signaling topologies? If copper solutions remain attractive, will I/O remain differential? This is great for a noisy backplane structures but if true coaxial routing were available would it still be the best choice? What would the speed requirements be? The public forum is reasonable starting point but history suggests that you will get your revenue return doing the same a technically similar task for customers who are implementing a proprietary solution.

Remember that for every hardware designer there must be 20 software engineers. This means that the burden of legacy hardware systems must pale in comparison to legacy control, provisioning and servicing software. I would guess that the legacy software would in many cases preclude adopting a fully standard hardware platform. This is because I assume that there will be requirements that will end up needing special functionality that just wasn't included in PICMG 3.0 for instance.

The basic question has been on the table for the past 10 – 15 years (see "No Time for Standards", April 1994 VMEbus Systems Magazine). Do competitors really want to switch customers to an industry standard box that their competitor may be able to fill and service more efficiently? Of course customers may desire to see this happen but what company would willingly do this for them. And would enough companies support this to make it a viable solution? We all officially say yes but only time will tell.

In the meantime, if the previous arguments are valid, it still makes sense to tackle the complex implementation of PICMG 3.0 because the premise that meshed architectures and meshed switches offer more flexibility and performance is correct. The challenges surmounted in implementing PICMG 3.0 will prepare a design team for any future differential serial fabric that is likely to be faced in the near future.