

Thoughts on the Tomahawk II

Russ White – <https://rule11.tech>

Best guess – written on October 2016. Russ should date his stuff.

Broadcom released some information about the new Tomahawk II chip last week in a press release. For those who follow hardware, there are some interesting points worth considering here.

First, the chip supports 256x25g SERDES. Each pair of 25G SERDES can be combined into a single 50g port, allowing the switch to support 128 50g ports. Sets of four SERDES can be combined into a single 100g port, allowing the switch to support 64 100g ports.

Second, there is some question about the table sizes in this new chip. The press release notes the chip has “Increased On-Chip Forwarding Databases,” but doesn’t give any precise information. Information from vendors who wrap sheet metal around the chipset to build a complete box don’t seem to be too forthcoming in their information about this aspect of the new chip, either. The Tomahawk line has long had issues with its nominal 100,000 forwarding table entry limit, particularly in large scale data center fabrics and applications such as IX fabrics. We’ll simply have to wait to find out more about this aspect of the new chip, it seems.

Third, there is some question about the forwarding buffers available on the chip. Again, the Tomahawk line has long been known for its very shallow buffers. While these generally aren’t a problem in well tuned data center fabrics (in fact, to the contrary, you don’t want a lot of buffering in time sensitive applications), there are situations where deeper buffers would be useful. The press release notes the new chip has “Higher Capacity Embedded Packet Buffer Memory,” but gives no details of what those larger buffers might look like.

Russ’ Take

While other vendors have supported 64x100g ports on their chips for the last year or so, this brings the 64x100g platform to a wider array of manufacturers, equaling many mid range (and potentially high end) switches available from mainline vendors. The fallout should be interesting in these areas; as Broadcom and other merchant silicon makers move up the performance scale, vendors are going to need to respond in some way. Analytics seems to be a natural path, but it won’t be long before chips in this area are also available from merchant silicon shops. Even this new chip, Broadcom says, has “Enhanced BroadView network telemetry features.” The value vendors add is being salami sliced into ever thinner pieces. This has implications for vendor partners, as well, of course—as the vendor’s added value drops through commodity components, there’s less “residual value” for a VAR to add.

It’s not all peaches and cream, though; there are still questions to ask, and things to think about. One issue with a switch operating at this density will be the faceplate; it’s difficult to see how the port counts on this chip will fit into a 1RU device (pizza box sized); manufacturers are going to need to be creative to build boxes in a smaller footprint that take full advantage of the potential here. The problems here aren’t just limited to physical real estate in the 1RU format, but also power consumption and heat dissipation, particularly in the optical units. It will be interesting to see what develops on this front.

This new chip looks interesting for a number of reasons, particularly for disaggregators and vendors who bundle the chip with their hardware and software to deliver a complete package. At the very least, merchant silicon is nipping at the tail of many vendor custom products—just another sign of how quickly the network market is changing.