Cover sponsor:
CHINDATA

capacity

Business intelligence for the global carrier industry

NEXT GENERATION HYPERSCALE
DATA CENTER

Capacity

ALEX JU

CHINDATA

IAN BAKAR.
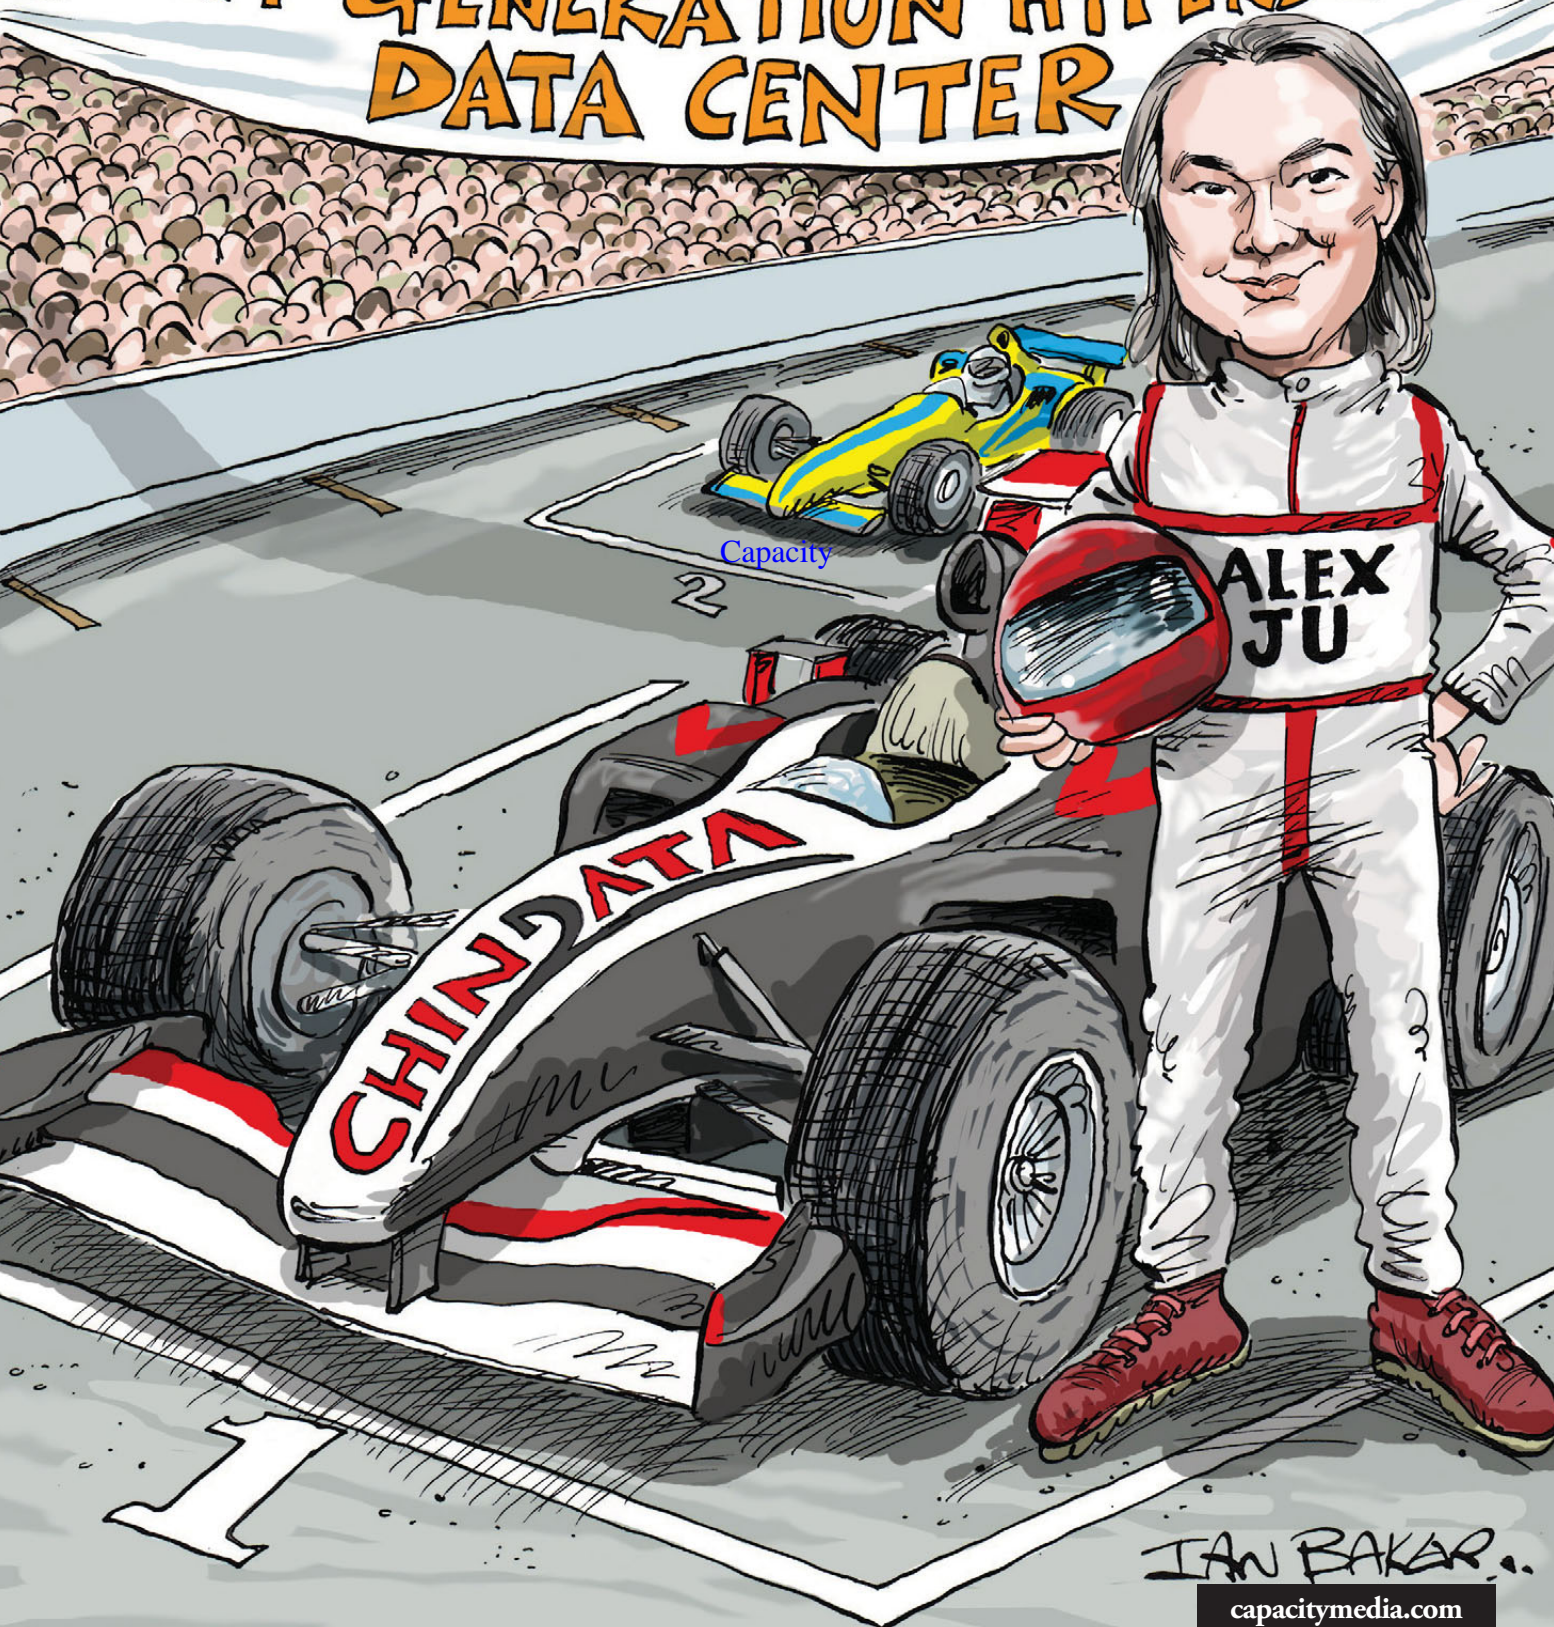
WHY CARRIERS ARE PULLING THEIR NETWORKS TO PIECES

Richard Brandon, vice president of marketing at RtBrick, tells Capacity why network disaggregation in the future of carrier networking.

We're all consuming more and more data over our internet connections, from online TV to 5G applications. We expect the network operators to deliver all of this traffic, but we don't want to pay any more for our broadband connections.

Consequently, the business of delivering internet services is getting harder and harder. Something has to change. Fortunately, it is about to.

Carriers are about to implement the biggest technology shift since the arrival of the internet — shaving at least half of the cost off their network procurement and operations. They refer to this shift as network disaggregation, which in simple terms means buying their network software separately from their hardware.

For many years the only way to deliver the levels of throughput required by the internet was to build the routing systems using custom silicon and optimise the software around it. That's why, in high-scale networking, software and hardware have always been intrinsically linked.

So, what's changed? In a word (well, two words) - merchant silicon. Silicon vendors now have the equivalent capabilities on their high-volume, low-cost networking chips that the traditional routing system vendors used to have in their customised systems. This merchant silicon is being used to build a new category of powerful low-cost 'bare-metal' switches, often constructed on the same outsourced assembly lines that manufacture the traditional router systems. These switches are a fraction of the cost of conventional telco routers, but just as powerful.

That's all very well for the hardware, but what about the software? Well, that's where a host of new companies are entering the market. They provide routing software that turns these bare-metal switches into IP/MPLS carrier routers, often specialising in different areas of the network, such as broadband access, edge and core. The approach that most of these new software providers is taking is very much like those used by the massive cloud-IT giants. Their software runs in a container on a Linux operating system on the switch and can be deployed with zero touch provisioning. It offers carriers the potential to turn their operations into something that resembles one of the 'internet-native' companies, rather than a traditional telco.

As well as reducing costs, the internet service providers now have the opportunity to mix and match hardware and software vendors. This makes it easier to upgrade the capacity of any dimension of the system without throwing away the existing infrastructure and will accelerate innovation for new services.

The next logical question, is whether all of this can be run in the cloud? Cloud infrastructure, mostly built with x86 servers, also comes with network cards based on this same merchant silicon. In fact, many telcos are using x86 servers to build out capacity for NFV (Network Functions Virtualisation). So, can they just use the same infrastructure for their network services?

Well, there are network topology and latency reasons why they might not want to send massive amounts of traffic in and out of centralised data centers, but that aside there is no reason they can't run the routing software on x86 servers. It's actually a good approach for small numbers of subscribers or for niche services with complex forwarding chains such as firewall functions (where you might run out of chip resources on a bare metal switch). But it doesn't scale as well for high volumes of standardised services, terminating millions of subscribers. In those environments, x86 servers can cost significantly more per port than bare-metal switches, especially for high speed ports at 100G and above.

But more problematic is the amount of power x86 servers consume. Every processing action has to be translated into x86 instructions – which is not an energy efficient way of forwarding packets. In fact, it uses about twenty times more power than a bare-metal switch.

So, for the bulk of high-volume internet services we can expect carriers to start deploying bare-metal-switches in large numbers. Taking advantage of the availability of new merchant silicon, and disaggregating hardware from software, promises to bring them some huge benefits:

- Much lower cost hardware – per port costs are a fraction of traditional systems
- Operators can swap and change software without throwing away or even upgrading the hardware
- It breaks the hardwiring between platforms and services, enabling easier re-use of common hardware across any service
- A bigger choice of software from a range of suppliers will make it faster to deliver new features and services
- Open standards and Web2.0 tools make the whole system easier to automate and to operate

These disaggregated network systems can replace many functions within a telco network, from core and edge routers to Border Network Gateways. And, along with a cloud-IT toolset to operate the networks, disaggregated systems are a fundamental part of the way telecoms operators are set to transform their operations to match the agility, simplicity and cost-levels of cloud-IT infrastructure.

So, now we can settle down to watch that online video stream without feeling quite so guilty about the amount of bandwidth we are consuming from our telco, in the knowledge that they have a new toolset to expand their network capacity at lower cost. It turns out that the future of the internet is all about pulling the network to pieces. ⌾