

REPORT REPRINT

Liquid goes multi-fabric with updated composable infrastructure platform

APRIL 22 2020

By John Abbott

Seven-year-old Liquid has moved beyond a reliance on using PCIe as a fabric and now supports multiple fabrics, including Ethernet. The product portfolio has been rebranded and expanded, making software the focus of its future strategy.

THIS REPORT, LICENSED TO LIQID, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, WAS PUBLISHED AS PART OF OUR SYNDICATED MARKET INSIGHT SUBSCRIPTION SERVICE. IT SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR RE-POSTED, IN WHOLE OR IN PART, BY THE RECIPIENT WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



Research®

Now a Part of

S&P Global Market Intelligence

Introduction

Composable infrastructure startup Liquid has broadened its positioning in the wake of a new round of funding at the end of 2019. Crucially, it's moved beyond a reliance on using PCIe as a fabric and now supports multiple fabrics, including Ethernet and InfiniBand. The product portfolio has been rebranded and expanded, making software the focus of its future strategy, with the Liquid Command Center software layer (purpose-built to support bare-metal composition) as the centerpiece. Liquid has also renewed its original technology partnerships and added significant new ones, including Dell, Intel, NVIDIA/Mellanox and Western Digital. A 20-GPU LQD8360 composable GPU system connected to Dell server nodes has recently been demonstrated for training large image classification models.

451 TAKE

The concept of composable infrastructure – making server, storage and networking infrastructure dynamically configurable on the fly – has been talked about for at least five years now, but so far its adoption has been limited. That's mainly because the underlying technologies that are required to support it continue to evolve at a rapid pace, making the end goal something of a moving target. Liquid sees a natural historical progression for on-premises architectures – from three-tier architectures to converged and hyperconverged, and now heading to composable and disaggregated. While we like the concept of composable infrastructure, we're still on the fence about just how widely applicable it will be within the spectrum of software-defined infrastructure, in general. However, it's undeniable that new data-intensive workloads, including artificial intelligence and edge computing, will increasingly rely on automation if they are to run efficiently while hiding some of the complexity. The current generation of static hardware can be limiting for these workloads without massive overprovisioning. Fabric independence is a significant step forward and enables a future path for customers (and for Liquid) as a new generation of fabric interconnects (such as CCIX, GenZ and CXL) capable of connecting and pooling high-bandwidth memory and accelerators from multiple sources starts to reach the market.

Context

Broomfield, Colorado-based Liquid was founded in 2013. It raised its first \$5.7m in seed funding in May 2015 from strategic investors Kingston Technology and Phison Electronics, alongside venture capitalists ABR Capital Management and DH Capital. The \$10m series A round in May 2017 was led by new investor Marker Hill Capital. Series B funding, closed at the end of 2019, added a further \$28m, taking total funding up to \$50m. This time the round was led by Panorama Point Partners, with participants Iron Gate Capital and DH Capital. The new money will be used to boost sales, marketing and engineering activities as the startup claims that its sequential revenue is doubling (from an undisclosed base amount) as the number of customers and deployments grow. Liquid says it's profitable. Original CEO and founder Jay Breakstone has moved on, but the other four cofounders are still in place. Sumit Puri (previously VP of marketing) has been in the CEO chair since April 2018. There are now 60 staff, up from 30 in 2017.

Product

Liquid's fabric-based composable infrastructure software disaggregates standard converged servers (with their own compute, storage and accelerators) into pools of sharable datacenter resources. The software delivers dynamic, bare-metal servers on demand, across PCIe Gen3 or Gen4 fabrics, as well as Ethernet and InfiniBand. The disaggregated hardware resources can include CPUs and networking, NVMe storage, GPUs and FPGA accelerators, and Intel Optane memory. These can be reconfigured and built out as required through the Liquid Command Center automation and orchestration software, with the aim of improving utilization rates and cutting down on the system footprint. Liquid Grid is the intelligent fabric switch, with up to 48 ports and 96 PCIe lanes, providing the low-latency interconnects between resource pools. Liquid also sells host bus adapters, expansion cabinets and the Element family of PCIe solid-state drive cards.

In partnership with Dell and Orange Silicon Valley, Liquid has put together the LQD8360 GPU server, linking a Dell server node at the bare metal layer with a separate JBOG (just a bunch of GPUs) physical enclosure with managed power and cooling, housing 20 NVIDIA Quadro GPUs in a 7U enclosure. Any number of GPUs can be dynamically assigned to any node on the fabric, with no physical chassis redesign required. Each GPU in the pod has 48GB VRAM, making 960GB VRAM in total. Single-, dual- or quad-node server configurations are available. Essentially, this creates a composable alternative to NVIDIA's DGX-2 GPU servers (which scale up to 16 GPUs), using standard server infrastructure. Other use cases include augmenting existing infrastructure (any host to any network) with new composable NVMe storage in an expansion chassis and the implementation of 'machine containers' (essentially templates for users wanting to define a machine for a specific application – perhaps SAP or Oracle, NVIDIA Cuda, VMware, Caffe or Tensorflow – that can be stored and launched very easily on demand).

Strategy

Liquid's business case is based on improving the utilization of expensive systems resources, so the recent interest in deploying GPU and FPGA accelerators for AI-related workloads only strengthens its story. Most customers struggle to bring their datacenter utilization rates beyond 12% on average – even Google manages only 30% on average. Liquid claims it can boost this up to 90% resource utilization. Software licensing also becomes more efficient, and different types of resources can be scaled independently, decoupling purchase decisions.

The multi-fabric capabilities are important – they bring in a wider range of resources (the combination of NVIDIA GPUs and Mellanox Ethernet adapters, for instance, or InfiniBand, which is widely used in the HPC sector). Further out, there are likely to be more fabric options as the market matures and new entrants emerge, such as GenZ and Intel's CXL (which depends on future versions of PCIe). Command Center provided unified management for the different types of fabric.

Key markets for Liquid are AI and deep learning (with GPU scale-out and sharing), dynamic cloud (i.e., the flexibility of cloud on-prem, but also for cloud service providers), HPC and clustering (massive scale-out and low-latency interconnect), visually intensive workloads and 5G/edge (scale-out and high-performance edge compute). The primary route to market will be through OEMs such as Dell, Inspur and Western Digital, as well as through channel partners, VARs and systems integrators.

Competition

Hyperconverged systems bring compute and storage resources together in a single virtualized platform on commodity hardware that's easy to manage and scale. Converged infrastructure packages more traditional enterprise compute, networking and storage components together as a turnkey, fully certified system. Neither is dynamic. Composable infrastructure promises the consolidation, disaggregation and pooling of resources in a much more dynamic way, making on-prem infrastructure or private clouds more cloud-like. Liquid sees the HCI and converged infrastructure vendors as the most direct competition, although some are partners at the same time.

HPE has been the most active participant in composable infrastructure with its Synergy platform, launched in 2016 and shipping a year later, and the OneView management platform that expands composability to a broader set of hardware platforms. HPE claims that revenue from Synergy's 3,000-plus customers topped \$1bn for the first time in 2019. Dell's alternative, the PowerEdge MX, was introduced in 2018. Both are more hardware-centric than Liquid and converged rather than disaggregated. Future architectures are likely to incorporate new high-speed fabric interconnect architectures like GenZ and Intel's CXL, which will help distribute compute, storage and memory resources, including GPUs and other accelerators. Those efforts are underway, but still not available.

Juniper's end-of-2018 acquisition of HTBase provided it with an infrastructure abstraction layer for compute storage and networking with its Juke multicloud container product, which it will use alongside the Juniper Contrail software-defined networking product.

Startups in the composable space include A3Cube, DriveScale, Gigalo and TidalScale. A3Cube (founded 2012) sells optimized appliances as the basis for its hardware-accelerated software-defined systems. DriveScale (founded 2012, funding \$26m) creates an elastic bare-metal cloud on top of any hardware that can deliver on-demand instances of compute, GPU and storage, including NVMe over Fabrics. Gigalo (founded 2012, funding \$4.5m) has introduced the FabreX composable network, which extends PCIe 4.0 to incorporate AI, big-data analytics and HPC. TidalScale (founded 2012, \$42m in funding) has focused on in-memory workloads (Oracle and SAP HANA) and big-data analytics.

SWOT Analysis

STRENGTHS

Liquid secured its latest funding before the current crisis and has some strong and established partnerships.

WEAKNESSES

The concept of composable infrastructure is five years on from its coining, and not yet established in the marketplace.

OPPORTUNITIES

AI, deep learning and other new workloads require the integration of new infrastructure to support them. Liquid can help with that.

THREATS

Evolving technology based on high-speed fabric interconnect architectures gives Liquid something of a moving target to aim at.