



**ANALYST REPORT**

# How Liquid is Dominating the Next Big Thing in the Data Center

**By:**

Arik Hesseldahl  
Analyst & Journalist





In any data center, there has for decades been one fundamental problem that's hard to get around: Business needs are fluid yet hardware is not.

Any time that the needs of a business change, and they change often, IT environments tend to be incredibly inflexible. Making changes can quickly spiral into expensive, time-consuming problems that can slow a business down, creating openings for competitors.

One reaction to this has been the "software defined" movements in both networking and storage. Resources that were once limited by the specifications of their hardware have been opened up by APIs and commodity hardware and made more flexible. But despite the potential, software defined networking and storage has made a limited impact on data centers overall.

But there's a new approach called composable infrastructure that has been gaining steam lately, and a company in

Colorado called Liquid has established itself as an early leader.

Put simply, composable infrastructure makes every part of the IT stack – computing, storage and networking and even GPUs – so flexible that it can be reconfigured on the fly as the needs change.

The economic implications by themselves are significant. IDC recently estimated that 75 percent of IT budgets are devoted to resources that are under-utilized. What this means is that in a typical enterprise environment under normal operating conditions, the vast majority of available computing capacity sits idle.

You wouldn't want to own a factory using less than a third of its capacity because owning equipment that's doing nothing is expensive. The same is true of IT infrastructure in a data center, and yet for a variety of reasons, utilization rates of 15 percent and less are the norm.



Companies like Google, Facebook and Amazon, whose data centers are often said to boast utilization rates are closer to 30 percent and sometimes higher are the exception to this rule. But they're not going to solve this problem for anyone but themselves. Liquid is promising to help companies reach target utilization rates of up to 90 percent.

How Liquid does this is interesting. It has created a PCIe switch that is placed atop a standard IT rack. That switch allows for the dynamic creation of servers configured at the bare metal level, where there's no virtualization software to add complexity or the added expense from licensing fees.

Once deployed Liquid allows IT resources to be pooled from as much or as little of the available hardware as needed in an arbitrary manner. Applications get the optimal level of resources they require, no more, no less. And if needs change, it's easy to re-configure on demand, and even in an automated way.

Let's say a workplace application gets especially busy around 10 AM on weekday mornings as everyone in the office signs on and gets to work. Liquid allows IT managers to set automated policies to spin up more resources — all the compute, networking and storage needed for a work day — from the available pool of hardware. Liquid enables the automatic reconfiguration of those resources once demand reaches a pre-set level, say 80

percent. And on weekends it can be set to devote a much smaller pool to that app, allowing the available hardware to be put to other uses, improving the overall utilization of the resources.

This young company — it has raised more than \$20 million in two rounds of VC funding — has a handful of proof-of-concept projects underway with large companies but has yet to name any customers. The potential use-cases are many. Big enterprise institutions which operate their own infrastructure could add Liquid switches to boost the efficiency of their applications and platforms, making them more competitive.

Emerging applications like artificial intelligence (AI), virtual reality (VR), and 3D rendering consume huge amounts of computing resources. Specialized GPU chips can speed up that process, but until now there's been no way to pool them. Liquid makes [that possible too](#), reducing a high compute workload that previously required days to one that now requires hours.

Overall Liquid is offering its customers more flexibility from their existing infrastructure than was possible with previous solutions. And when they invest in new hardware, they can either get more computing bang for the buck, or scale back their budgets and simply not buy extra gear they don't need, essentially enabling them to do more with less.



## WHY LIQID IS THE LEADER IN COMPOSABLE

Liquid also has some important wind at its back. Earlier this year semiconductor giant Intel released version 2.1 of its Rack Scale Design (RSD) [specification](#), basically a guide to what the Intel thinks data center infrastructure will look like in the near future. Liquid's approach to composable is very much aligned with the promise of Intel's RSD.

A key plank of that specification includes the use of PCI Express as the preferred interface for connecting all pools of resources together. There's a handful of reasons that PCIe is a better choice than Ethernet, which is what a competing composable infrastructure solution called Synergy from Hewlett Packard Enterprise favors. For one thing Ethernet is slower and tends to suffer from higher latency than PCIe. Additionally, Ethernet is not a native interface for connecting to GPUs or to SSD storage, which PCIe is.

Liquid, it turns out, is the only composable infrastructure player in the marketplace today that can unite all the different elements of the stack – compute, GPU, storage and networking – through this common PCIe connection.

It did this by making a PCIe fabric switch to interconnect the pools of resources, because the interface enables all the major datacenter elements and peripherals to communicate natively. Liquid's management console software keeps track of it all.

Amazon, whose data centers are often said to boast utilization rates are closer to 30 percent and sometimes higher are the exception



# DISAGGREGATED VERSUS CONVERGED

Another key aspect of Liquid's approach is that it enables a disaggregated approach to hardware. This means that instead of a single piece of hardware that combines compute, memory and storage resources, it offers the ability to connect standalone compute box, storage box, or GPU box to scale out only the required resource.

The company believes that the disaggregated approach is the right one for one key reason: Flexibility. Customers only have to pay for resources they need and can forego the ones they don't. This approach differs from what other players in the composable space offer. Most favor the converged approach of selling boxes that combine all the datacenter elements, whether or not they are necessary to what a given user is trying to achieve.

What this effectively means is that you can buy precisely the mix of compute, storage, and GPUs that meets your needs, without having to pay for additional elements that you might not want or need.

GPUs in particular are the new cool kid in the data center, adding some extra computing muscle to specialized work that requires more time and energy with traditional CPUs. Liquid's ability to harness hundreds of them into a single cluster is unique among the composable infrastructure players.

Down the road, I think the most exciting direction composable infrastructure may take lies with its potential for machine learning and artificial intelligence. Already Liquid's management tools allow for the use of policies and rules that trigger automated responses under certain conditions.

Eventually systems will learn from patterns of usage and demand and configure themselves in ways that can't help but evoke philosophical discussions that include references to science fiction movies, but which will in day-to-day reality save time, energy and the human effort associated with the massive undertaking that managing a data center already is.