

Technical Validation

# Liquid Composable Disaggregated Infrastructure

## Configuring Servers On-demand in the On-premises Data Center

By Alex Arcilla, Senior Validation Analyst

April 2021

This ESG Technical Validation was commissioned by Liquid and is distributed under license from ESG.



## Contents

Introduction ..... 3

    Background ..... 3

Liquid Composable Disaggregated Infrastructure ..... 4

ESG Technical Validation..... 5

    Accelerated Time to Value..... 5

    Increased Data Center Resource Efficiency ..... 7

    Improved IT Agility ..... 9

The Bigger Truth..... 11

### ESG Technical Validations

The goal of ESG Technical Validations is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Technical Validations are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team’s expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.

## Introduction


This ESG Validation Report documents testing and evaluation of Liquid Composable Disaggregated Infrastructure (CDI), with a goal of validating how organizations can accelerate time to value, increase resource efficiency, and improve IT agility within their on-premises data centers.

## Background

Organizations continually seek ways to modernize how they can provision computing and storage capacity in their data centers for workloads quickly and flexibly in order to improve business agility. To that end, ESG research on data center modernization investment has uncovered that 34% of organizations anticipate increasing their use of on-premises hyperscaler cloud solutions in the next 12-18 months,<sup>1</sup> while 32% anticipate a significant investment in implementing a software-defined data center strategy (see Figure 1).<sup>2</sup>

**Figure 1. Top Six Areas of Data Center Modernization for Investments in the Next 12-18 Months**

**In which of the following areas of data center modernization will your organization make the most significant investments over the next 12-18 months? (Percent of respondents, N=664, five responses accepted)**



Source: Enterprise Strategy Group

Organizations have relied on manually provisioning, installing, and scaling server hardware to support a wide range of workload requirements. Yet, because the time to quote, ship, and provision these servers typically ranges from weeks to months, the time to extract the business value from the workloads ends up being protracted and costly.

In recent years, organizations have turned to infrastructure-as-a-service (IaaS) to deploy and scale virtualized servers on-demand, both in the public cloud and on-premises (via hyperscaler cloud solutions). While gaining advantages of speed, agility, and scalability, sizing workloads can be difficult when using cloud compute instance types with predefined amounts of CPU or GPU capacity, RAM, and storage. The risk of over- or under-provisioning to right size enough instances for a workload is always present. Those opting for hyperscaler cloud solutions must also purchase hardware that is designed to run these compute instances from cloud service providers (CSPs). Either option also poses the risk of incurring unwanted cloud-related expenses.

<sup>1</sup> A hyperscaler cloud solution is a hardware and software stack offered by a cloud service provider (CSP), such as Amazon Web Services (AWS) Outposts, that enables businesses to deploy and run cloud-like infrastructure (such as compute instances) spanning on-premises data centers and public clouds using a common set of tools and operating procedures.

<sup>2</sup> Source: ESG Master Survey Results, [2021 Technology Spending Intentions Survey](#), December 2020.

To address these challenges, what if organizations had a solution that helps to configure and scale servers on-demand, via software, using disaggregated components already located on-premises? Once a workload is no longer needed or has changed, what if that platform can repurpose those same components to new workloads, eliminating the need to procure new servers or upgrade servers manually, thus not incurring additional capital and operational expenses?

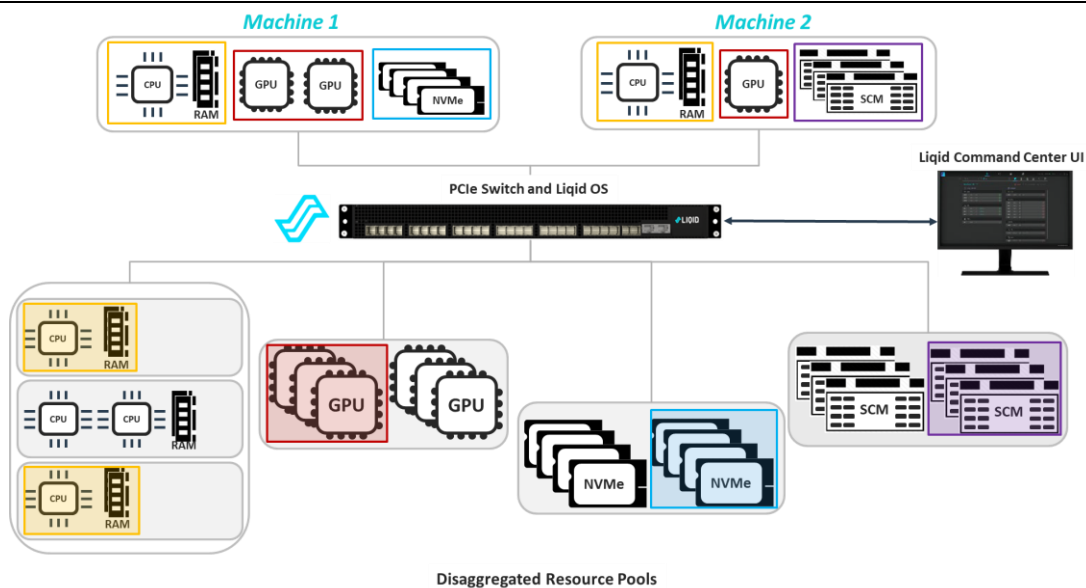
### Liquid Composable Disaggregated Infrastructure

Liquid Composable Disaggregated Infrastructure (CDI) is a software-based solution that can enable organizations to provision bare-metal servers on demand. The solution brings the flexibility and agility of IaaS to on-premises data centers by enabling organizations to quickly configure servers using pools of hardware resources, including compute,<sup>3</sup> GPUs (e.g., NVIDIA A100), NVMe SSDs, storage-class memory (SCM)<sup>4</sup>, FPGAs, and NICs. With Liquid CDI, organizations can quickly add or subtract GPU, accelerator, storage, and/or networking resources from existing servers via software, eliminating the need to physically install or remove components from a server chassis. This allows organizations to deploy and scale the exact quantity and type of resources new workloads require.

Figure 2 illustrates how servers are composed with Liquid CDI over a PCIe fabric. Each pool of disaggregated resources is interconnected via a PCIe Gen3 or Gen4 fabric switch. Host servers must be installed with a PCIe HBA. Liquid Matrix Software, installed on the switch fabric, enables organizations to configure server with GPUs, NVMe SSDs, and SCM without having to manually connect the devices to each other via Liquid's UI, API, or CLI. No Liquid-specific drivers or agents are needed for the server to recognize these added devices. Resources can be added to running systems without a reboot.

Should the workload require additional GPU processing or storage capacity, the required components can simply be added via software. All components can be scaled independently of each other to meet any workload's requirements. If the workload is no longer running or requires fewer resources, those components can then be returned to the free pool for relocation to other servers. These software-defined server configurations act as bare-metal servers that support common Oses, hypervisors, or container engines.

**Figure 2. Liquid Composable Disaggregated Infrastructure**



Source: Enterprise Strategy Group

<sup>3</sup> CPUs and RAM are not disaggregated in Liquid CDI. Servers already contain a defined amount of CPU and RAM capacity.

<sup>4</sup> With Intel Optane technology, SCM components can extend system memory.


Liquid CDI also supports both Ethernet and InfiniBand (IB) fabrics, providing organizations flexible connectivity options, including NVMe-over Fabric (NVMe-oF). With high-speed Ethernet switches, organizations can leverage Liquid CDI to compose servers with the required amount of NVMe SSD storage capacity using SmartNICs.<sup>5</sup> By leveraging Liquid’s support for multiple switch fabric types, organizations can compose a single host server with resources interconnected with both PCIe and Ethernet/IB fabrics.

## ESG Technical Validation

ESG performed evaluation and testing of Liquid CDI at the company’s facilities in Broomfield, CO. Testing was designed to demonstrate how the solution can help organizations accelerate time to value, improve IT agility, and increase resource efficiency.

ESG began with the test bed shown in Figure 3. One Liquid 48-port PCIe Gen4 Fabric Switch interconnected with two compute modules within an Intel server system S9200WK equipped with Intel Xeon Processors. Each contained a Liquid PCIe Gen4 x16 Fabric HBA card. The Liquid Fabric Switch also connected to two Liquid expansion chassis, each containing eight PCIe Gen4-enabled slots. Both chassis were populated with a total of 16 NVIDIA A100 Tensor Core GPUs. The Liquid Director interconnected the Liquid Fabric Switch and expansion chassis to enable server composability.

**Figure 3. Liquid Test Bed for Composing a New Server**



Source: Enterprise Strategy Group

## Accelerated Time to Value

Configuring a server for a workload in minimal time is critical in today’s business climate to meet business demands quickly and efficiently. However, the traditional process of provisioning and configuring server hardware (determining workload requirements; purchasing the required hardware components; then assembling, testing, and deploying the server) can take weeks or months to complete. With Liquid CDI, organizations can drastically simplify this process by configuring the servers on-demand via software, thus accelerating the time to meet business demands.

## ESG Testing

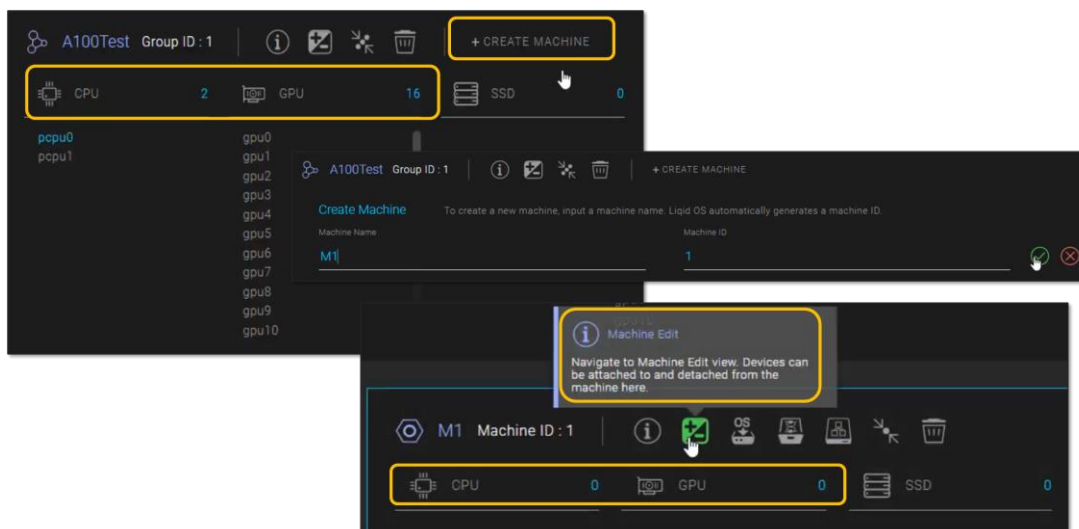
ESG first configured a server requiring a large amount of GPU compute capacity to run a hypothetical artificial intelligence/machine learning (AI/ML) workload. We viewed the available hardware resources for configuring servers via

<sup>5</sup> Liquid CDI also currently supports InfiniBand switch fabrics.

the Command Center UI (see Figure 4), including two CPUs from our test bed, “pcpu0” and “pcpu1,” and the 16 NVIDIA A100 GPUs, labeled “gpu0” through “gpu15.” All components were assigned to a group<sup>6</sup> called **A100Test**.

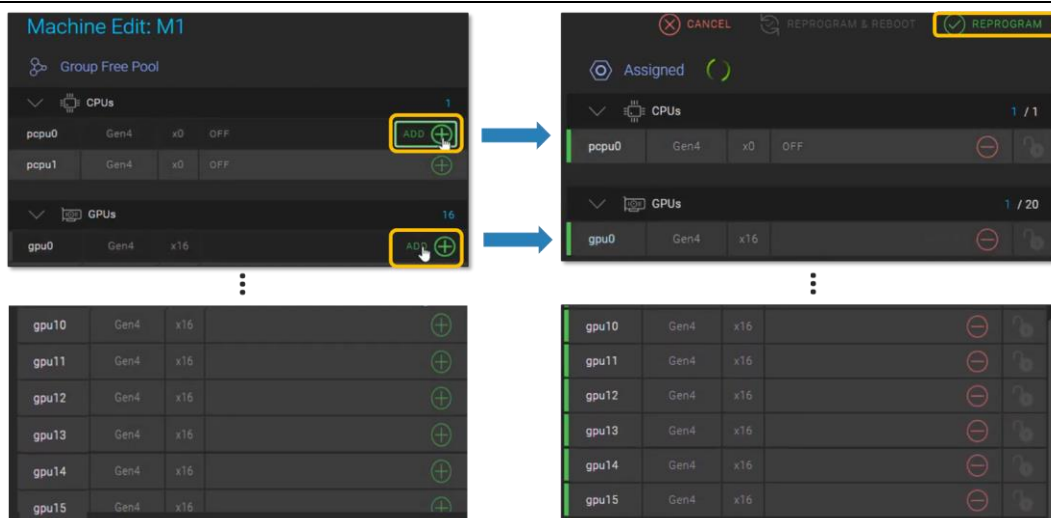
We began by clicking on **Create Machine** to compose an Ubuntu server and named it “M1.” We noted that “M1” had no CPUs or GPUs assigned to it. We then clicked on the **Machine Edit** icon to assign resources to “M1.”

**Figure 4. Creating a New Machine “M1”**



ESG then assigned “pcp0” and all 16 GPUs to the “M1” machine by clicking **Add** next to each line item (see Figure 5). Recently added components were flagged by the green vertical line next to each component. After “pcp0” and the 16 GPUs appeared under the **Assigned** column, we clicked on **Reprogram** in the upper right-hand corner to compose “M1.”

**Figure 5. Adding Compute and GPU Resources to Machine M1**



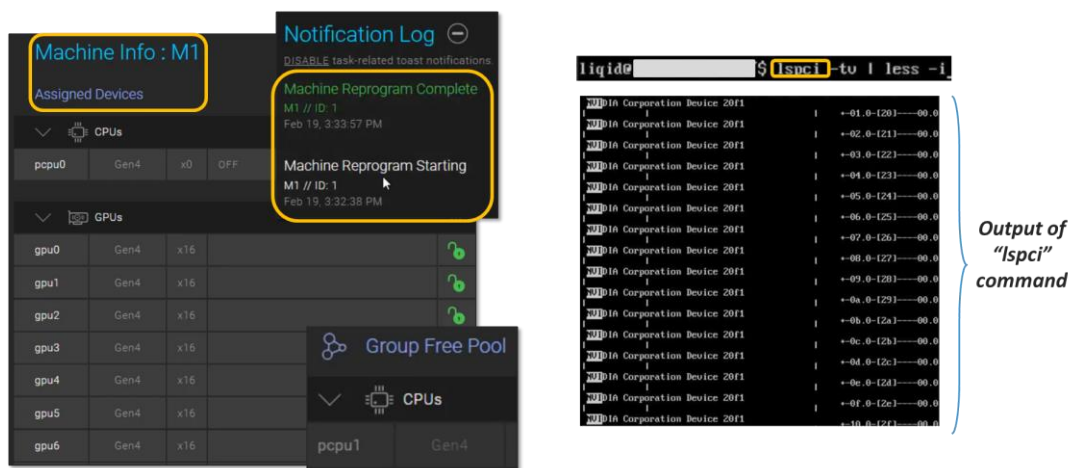
ESG observed that the **Reprogram** step completed in under five minutes via notifications (see Figure 6). We also saw that these components were in the **Assigned** category when viewing the “M1” machine. In other words, “M1” was a bare-metal

<sup>6</sup> A group enables an IT administrator to allocate specific resources to a user or team. The group can be restricted with Liquid's role-based access control (RBAC) functionality.

Ubuntu server connected electrically via PCIe Gen4 to the 16 GPUs, as if we physically installed 16 GPUs into a server. We also noted that the only resource available for configuration was “pcpu1,” the second Intel Ubuntu server.

To confirm that the GPUs were connected via PCIe to “M1,” ESG remotely accessed it and ran the “lspci”<sup>7</sup> utility. All 16 GPUs were listed as attached to “M1” and ready to be utilized.

**Figure 6. Verifying that Available GPUs are Connected to “M1” Machine**



Observing the relative ease of configuring servers on-demand, ESG noted the value of Liqid CDI goes beyond its ability to configure servers on-premises for any given workload in real time. By drastically shortening the IT provisioning cycle, organizations can immediately utilize data center resources and increase their return on investment (ROI) in less time.

### Why This Matters

Configuring servers to support a workload’s requirements is tedious and time consuming when using traditional IT provisioning processes. Ordering the proper hardware and additional components, configuring hardware, testing, and deploying new servers can last a few weeks to a few months. With this approach, organizations cannot meet their critical business needs in a timely manner.

ESG validated that Liqid CDI accelerates time to value by enabling organizations to build and configure servers to meet precise resource requirements for any given workload in a drastically reduced amount of time compared to common IT provisioning cycles. We noted that a significantly shortened server provisioning process can help organizations to extract maximum business value from newly deployed servers in a minimal amount of time. Subsequently, deploying servers to run workloads almost immediately helps organizations to meet business demands quickly and efficiently, without having to schedule maintenance windows or refreshes.

### Increased Data Center Resource Efficiency

Organizations attempt to right-size servers for supporting and scaling new workloads. However, servers are overprovisioned at the onset to accommodate future growth, and resources are underutilized. When organizations scale or tear down workloads, a surplus of GPU compute and storage capacity can accrue. Reusing and redeploying this surplus is not an easy exercise. While organizations may attempt to “shoehorn” a workload onto a pre-configured server, that server either does not have sufficient compute or storage capacity to support the new workload’s requirements, or the server has


<sup>7</sup> “lspci” is a command on Unix-like operating systems that lists detailed information about all PCI buses and devices in a given system.

excess compute or storage capacity that will be wasted on a smaller, less demanding workload. With Lijid CDI, organizations can maximize utilization of data center resources in real-time, freeing them up for other workloads.

### ESG Testing

ESG modified the test bed shown in Figure 3. We used a 48-port Lijid PCIe Gen3-enabled switch (versus Gen4). We replaced the two Intel-based Ubuntu servers with one Dell EMC PowerEdge 240 and one Dell EMC PowerEdge 640 server, each running on Intel Xeon processors and containing a PCIe Gen3 host bus adapter (HBA). We also replaced the two Lijid PCIe Gen4-enabled expansion chassis with two Lijid expansion chassis containing a total of eight PCIe Gen3 enabled slots. One chassis housed two NVIDIA V100 GPUs, 1.5TB Lijid Intel Optane capacity, and 1.6TB Lijid NVMe SSD capacity, while the other contained two NVIDIA T4 GPUs and 3TB Lijid NVMe SSD capacity (see Figure 7).

**Figure 7. Modified Test Bed with Dell EMC Servers**

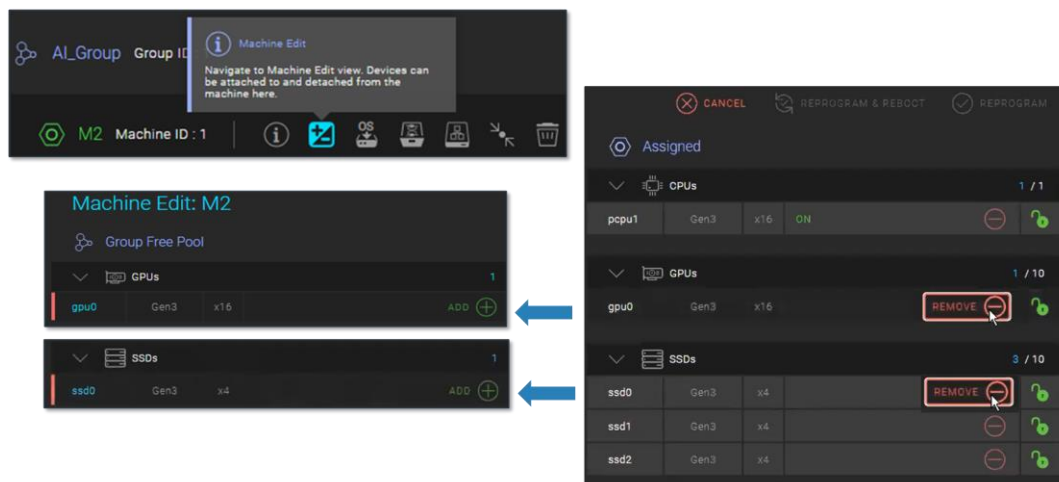


Source: Enterprise Strategy Group

ESG proceeded to remove resources from the “M2” machine, a composed server consisting of the Dell EMC PowerEdge 640 containing the CPU named “pcpu1,” one NVIDIA V100 GPU, and three Lijid NVMe SSDs (see Figure 8). We clicked on the **Machine Edit** button and saw these devices associated with “M2.” We scaled the server down by one GPU and one SSD by clicking on **Remove** next to “gpu0” and “ssd0.” These devices were automatically reassigned to the **Group Free Pool**.



**Figure 8. Removing Excess GPU Compute and Storage Capacity from “M2” Machine**



In this case, ESG saw how organizations can quickly free up resources from servers, such as when the workload no longer requires its original amount of GPU compute and storage capacity. Rather than allow the excess capacity to remain idle, Liquid CDI enabled us to remove those components for other workloads. For example, organizations can maximize GPU resource utilization during off hours by temporarily reallocating GPU resources from a VDI workload running on “M2” to workloads that run during off hours.

### **Why This Matters**

Maximizing data center resource utilization is challenging when scaling or tearing down workloads. Existing compute and storage resources become available, yet reusing them is not straightforward. Either the excess GPU compute and storage capacity is wasted on a workload with less stringent requirements, or the capacity remains idle until moved to workloads that need it.

ESG validated that the Liquid CDI enables organizations to maximize data center resource utilization. By observing how simply we could remove Nvidia V100 GPUs and Liquid SSDs from an existing server configuration, we could see how GPU compute and storage resources are made available to be assigned to other workloads. Organizations can minimize capital expenses since these available resources can be reused by other workloads when needed.

### **Improved IT Agility**

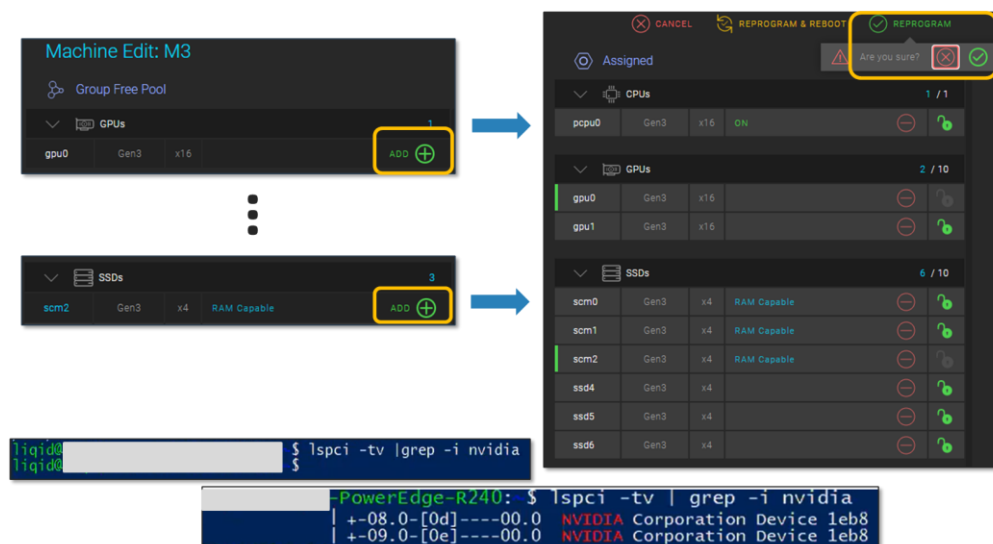
Modifying and scaling existing server infrastructure to accommodate changes in workload requirements incurs downtime, preventing any business from meeting demands in a timely manner. Organizations must navigate the IT provisioning process to either install additional components to an existing server or build additional servers to handle an increased load. Satisfying modified workload requirements will again require a few weeks to months, depending on the extent of the hardware changes to be made. On the other hand, Liquid CDI can help organizations to dynamically modify composed servers in real time, without the need to physically modify existing server infrastructure.

### **ESG Testing**

Using the same test bed in Figure 7, ESG scaled up GPU compute capacity using resources reclaimed from “M2” and added SCM (Intel Optane memory capacity from the **Group Free Pool** to an existing machine named “M3.” We navigated to the **Machine Edit** page associated with “M3” (see Figure 9) and viewed that it was comprised of the Dell EMC PowerEdge 240

("pcpu0") with one Tesla V100 GPU ("gpu1"), three Liquid PCIe Gen3 SSDs ("ssd4" through "ssd6"), and two Liquid PCIe Gen3 SCM devices ("scm0 and "scm1").

**Figure 9. Scaling Up GPU Compute and Storage Capacity on Existing "M3" Machine**



ESG then reassigned "gpu1" and added "scm2" to "M3" by clicking on the **Add** signs for these two line items under the **Group Free Pool**, as indicated by the green vertical lines. We then clicked **Reprogram** so that they would be connected via PCIe to "M3." Once this task completed, we ran the "lspci" utility on the Dell EMC PowerEdge 240 server and saw both GPUs listed. We also verified the addition of "scm2" with "lspci."

ESG took note that the use of Liquid CDI not only eases the process of scaling both GPU compute and storage capacity within an existing server but also increases the chances that the updated server configuration is implemented correctly. Installing additional components into any server chassis is not difficult, but the manual work required can consume time unnecessarily and introduce unanticipated issues that prevent a successful server reconfiguration. With Liquid CDI, such errors are practically eliminated.

### **i Why This Matters**

To remain agile and responsive to business needs, organizations will need to scale certain workloads in real time, such as an AI/ML workload that must deliver real-time insights, even as the amount of data grows. However, scaling GPU compute and storage capacity in real time on-premises is simply not doable with today's IT operational processes.

ESG validated that the Liquid CDI helps organizations to modify and scale both GPU compute and storage capacity of existing on-premises servers in a matter of minutes with Liquid Command Center. We noted the simplicity of attaching additional Nvidia V100 GPUs and Liquid SSDs to a server in order to accommodate scaled-up workload requirements. We accomplished these tasks remotely without opening up an existing server chassis and manually installing components, helping to minimize operational expenses.

## The Bigger Truth

Extracting the maximum value out of a server infrastructure in the data center is difficult when workload requirements change. While traditional IT provisioning processes enable organizations to meet workload requirements exactly, the risk of incurring unnecessary capital and operational expenses is high. If the workload scales up, more compute or storage capacity needs to be purchased and installed. If the workload scales down, excess capacity remains idle until another workload, with slightly mismatched GPU compute and storage requirements, can be supported. In both cases, capital and operational expenses increase over the long run, the ROI decreases, and the business cannot meet business demands in an efficient and agile way.

Liquid CDI enables organizations to compose servers on demand within on-premises data centers without wasting hardware resources. Leveraging pools of disaggregated GPU compute and storage devices, organizations can compose servers via software to meet exactly any workload's requirements. Once requirements change, Liquid CDI helps to modify server configurations with minimal delay. Should the workload be scaled back or torn down, components of the composed server are reclaimed by the resource pools, making them available for new workloads to reuse them. The average IT provisioning cycle can now last seconds or minutes, not weeks or months (assuming that the organizations have the proper number of components in the free pool to support new and existing workloads).

Throughout our testing and evaluation, ESG validated that Liquid CDI can help organizations to compose servers on the fly, remove excess GPU compute and storage capacity from existing server configurations, and reassign those same components to another workload in a matter of minutes using Command Center. We observed that we could add any number of device types (GPUs, NVMe SSDs, and SCM) to any existing server equipped with CPU and RAM, as each device can be scaled independently of each other. We verified that the devices were indeed attached to the server configurations we created using LINUX-based utilities.

While Liquid CDI can improve how organizations build out and manage server infrastructure on-premises, it will be important to evaluate how the solution will include other functionality to support these tasks. Those that come to mind include integration with additional automation and orchestration tools and policy-based management via templates to configure multiple servers based on workload type. However, ESG acknowledges that the roadmap evolution is only a matter of time.

ESG was impressed with the software-based solution Liquid offers to help organizations rethink their IT provisioning processes and uncover value from an on-premises data center infrastructure. We highly recommend taking a closer look at Liquid CDI should you want to achieve that end.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



**Enterprise Strategy Group** is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.



[www.esg-global.com](http://www.esg-global.com)



[contact@esg-global.com](mailto:contact@esg-global.com)



508.482.0188