

A Data for Good Solution empowered by VMware Cloud Foundation with Tanzu



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Table of Contents

1	Introduction	3
2	Technology as a force for good	3
2.1	Data for Good Initiative:.....	4
3	Modernized Infrastructure with VMware Cloud Foundation	4
3.1	Demands of Modern Applications.....	4
3.2	VMware Cloud Foundation with Tanzu	5
3.3	GPUs for Machine Learning	6
3.4	vSphere 7 with Bitfusion brings GPU sharing capabilities.....	7
4	Machine Learning Platform for Data Scientists:	7
4.1	Benefits.....	8
4.2	Features	8
4.3	The working of vMLP	9
5	Solution Infrastructure:	10
5.1	VMware Tanzu Kubernetes Infrastructure:.....	10
5.2	GPU Cluster:	11
6	Solution in Action	12
7	Conclusion:	15



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

1 Introduction

Nonprofits typically work to create awareness about the rights of the distressed and deserted sections of the society. Quite often unfortunately due to lack of resources and proper guidance, they fail to reach their true potential.

The modern world creates and has access to large swaths of data, which can be potentially harnessed to gain insights. Data is one of their critical assets, which when combined with technology can provide critical and valuable insights. Nonprofits across the world are slowly realizing the power of data to improve their operational efficiency, cost management, donor outreach, and planning for the future.

There is significant potential for the usage of data analytics combined with machine learning to harness the power of data to benefit nonprofits. Data analytics is a critical tool that can be used for discovering, diagnosing and deploying models that can be used by nonprofits to better serve their community. There are many examples of nonprofits that are usually averse to data-driven decisions leveraging data for good use.

2 Technology as a force for good

During the 2019 VMworld, VMware CEO Pat Gelsinger asked the question: Is tech good or bad?," "The answer is generally neutral; it's neither good or bad, but it's often how we shape it and will technology shape the world that we want to live in, or will it create a world that we're afraid to live in? I sincerely and equally believe that technology has the opportunity to expand the life of every human on the planet and eradicate diseases that have plagued mankind ... to give modern education to every child on the planet ... to lift the remaining 10% of the planet out of poverty, reverse the implications of climate change," he continued.



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Technology holds a lot of potential to be a force for good. Data is one of the key elements that can be combined with technology as a force for good.

2.1 Data for Good Initiative:

The data for good initiative was inspired by Pat Gelsinger (VMware CEO) “Tech as a force for Good”. The goal was to leverage the VMware Cloud Solutions lab resources for Good. Can we build a solution that leverages VMware Cloud Foundation, Tanzu and the modern hardware available for the VCSLAB to serve the technology needs of nonprofits?

3 Modernized Infrastructure with VMware Cloud Foundation

3.1 Demands of Modern Applications

Applications are essential...and modern applications are scaling at rates never experienced before. IDC predicts that over 500 million apps using cloud native methodologies over next 5 years, which is more than was developed in past 40 year combined.

But the applications are become much more complex and diverse. As a result, application innovation is driving a need for infrastructure modernization. Modern apps require infrastructures that support all of their components. Without an adaptive infrastructure, developers and IT operators are often at odds and struggle to provide the services that apps require.



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

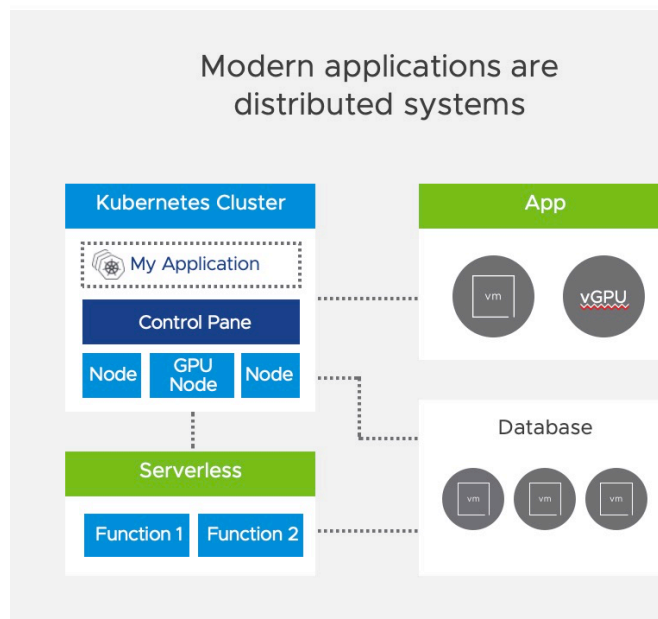


Figure 1: Distributed nature of modern applications

Developers essentially expect the infrastructure to just work with their cloud native applications. IT admins are using software-defined technologies to modernize their infrastructure to make it more agile for cloud native developers, while ensuring security, stability and consistency across environments and across clouds, while minimizing their security exposure. This is where many organizations struggle – how can you meet the needs of the fast-paced developers and DevOps teams while ensuring that infrastructure can be agile, secure and stable?

3.2 VMware Cloud Foundation with Tanzu

VMware Cloud Foundation with Tanzu delivers hyper-speed Kubernetes that provides agility, flexibility and security for modern apps. VMware Tanzu delivers the infrastructure, and services to meet changing business needs to rapidly deploy new applications. VCF provides consistent infrastructure and operations with cloud agility, scale and simplicity.



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

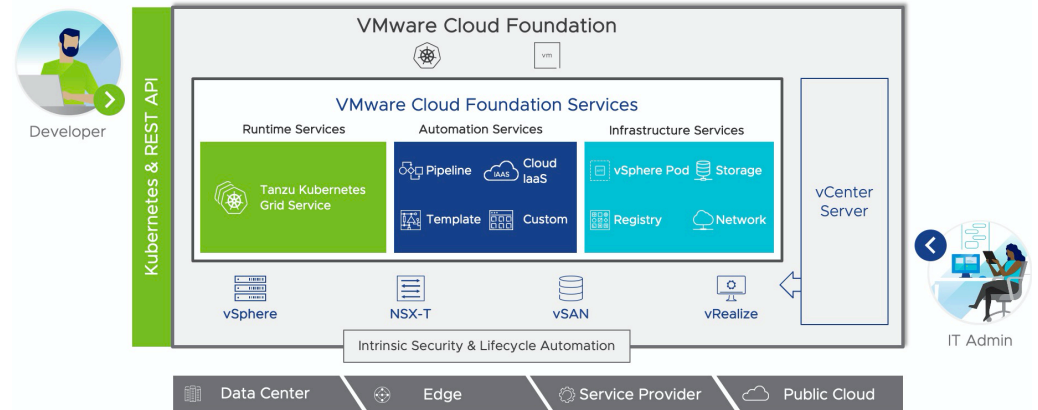


Figure 2: VMware Cloud Foundation with Tanzu

VMware Cloud Foundation with Tanzu is a Hybrid Cloud Platform that accelerates development of modern applications that automates the deployment and lifecycle management of complex Kubernetes environments.

- IT admins have complete visibility and control of virtualized compute, network and storage infrastructure resources through VCF.
- Software defined compute, storage and networking with vSphere, NSX-T and vSAN/VVOL provides ease of deployment and automation.
- Developers have frictionless access to Kubernetes environments and infrastructure resources through VCF Services.
- VMware Cloud Foundation provides runtime services automation services and infrastructure Services, all delivered via Kubernetes and RESTful APIs

3.3 GPUs for Machine Learning

With the impending end to Moore's law, the spark that is fueling the current revolution in deep learning is having enough compute horsepower to train neural-network based models in a reasonable amount of time. The needed compute horsepower is derived largely from GPUs, which NVIDIA began optimizing for deep



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

learning since 2012. A lot of the machine learning and AI related work relies on processing large blocks of data, which makes GPUs a good fit for ML tasks. Most of the machine learning frameworks have in built support for GPUs. There is a need to provide the capabilities needed by data scientists such as GPU access from Kubernetes environments.

3.4 vSphere 7 with Bitfusion brings GPU sharing capabilities

With the release of vSphere 7 many a new feature with vSphere Bitfusion was introduced. vSphere Bitfusion provides the ability to share NVIDIA GPUs over the network. Modern applications quite often need access to GPUs and their massive compute capabilities for timely and efficient processing. There is a significant need for Kubernetes based developer environments to access GPUs. A combination of the VMware Tanzu platform and vSphere Bitfusion for GPU sharing over the network can help meet the infrastructure needs of modern data scientists.

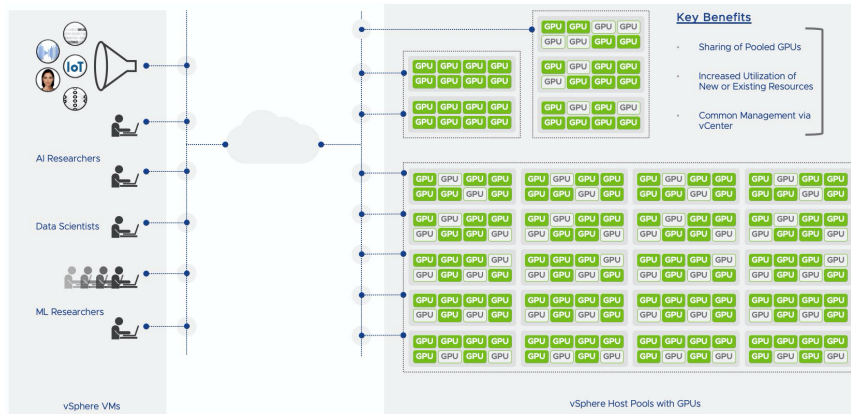


Figure 3: Leveraging vSphere Bitfusion for GPU access over network

4 Machine Learning Platform for Data Scientists:

The data for good program needs a platform for data scientists to be able to import and analyze data for non-profits. The data scientists



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

will use this platform to gain insights from critical data from non-profits to help them better server their base.

vMLP is a platform to provide an end-to-end ML platform which seamlessly deploys and runs AI/ML workloads on top of VMware infrastructure.

4.1 Benefits

The vMLP platform provides the following benefits:

- Enables cost efficient use of shared GPUs for AI/ML workloads
- Reduces the risks of broken Data Science workflows by leveraging well-tested and ready-to-use project templates
- Faster development for AI/ML models by utilizing end-to-end tooling including fast and easy model deployment served via standardized REST API

4.2 Features

The vMLP platform has the following features:

- Deployment of familiar Jupyter Notebook workbenches
- JupyterLab extensions to achieve smooth workflows
- Support for distributed model training leveraging popular OSS frameworks and GPUs
- Storing and reliably tracking experiments for reproducibility
- Efficient turn-around on model testing based on the built-in serving framework
- Multi-cluster Federated ML support
- Sharing metadata about Data Sources to enable team collaboration



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

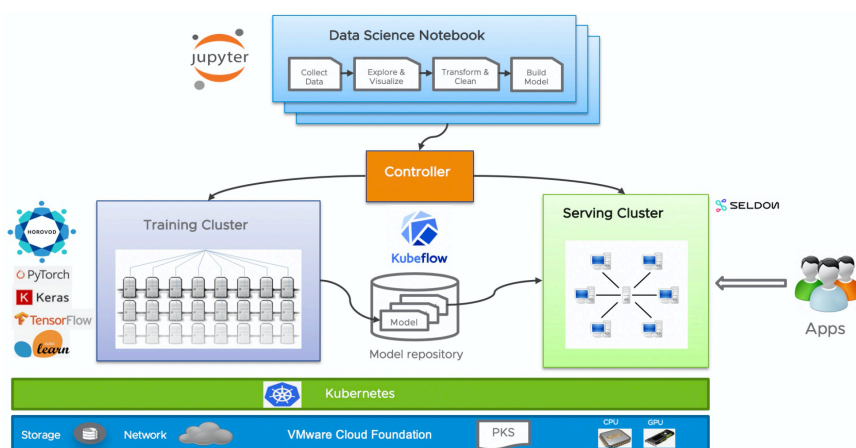


Figure 4: The vMLP platform components

4.3 The working of vMLP

vMLP is based on Kubeflow which can be compared to a Linux kernel. Kubeflow is not consumed directly, but through some distros like Ubuntu, which ensure more thorough testing, integration with diverse drivers, usability fixes, etc. vMLP is a Kubeflow distro with an enterprise and usability focus available as a [VMware fling](#).

Status	Name	Age	Image	CPU	Memory	Volumes
✓	a2	21 hours ago	vmkp-tensorflow-2.1.0-notebook-cpu-0.3.1-alpha2	0.5	1.0Gi	CONNECT
✓	afadfa	9 days ago	vmkp-tensorflow-2.1.0-notebook-cpu-0.3.1-alpha1	0.5	1.0Gi	CONNECT
✓	t1	1 day ago	vmkp-tensorflow-2.1.0-notebook-cpu-0.3.1-alpha2-test1	0.5	1.0Gi	CONNECT
✓	t2	1 day ago	vmkp-tensorflow-2.1.0-notebook-cpu-0.3.1-alpha2-test2	0.5	1.0Gi	CONNECT

Figure 5: Jupyter Notebook Servers offered with vMLP

vMLP can leverage enterprise-grade VCF components: etc. to provide Kubeflow significant accelerated infrastructure. vMLP also includes additional services which are not part of Kubeflow such as MLflow Model Repository, BentoML, Greenplum, R Kernel, etc.



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

5 Solution Infrastructure:

The Logical schematic of the infrastructure used in the solution is shown below. A Tanzu Kubernetes based compute cluster provides scalable compute infrastructure and a vSphere Bitfusion based GPU cluster provides required. A high speed low latency network is leveraged by Kubernetes workloads to access GPU resources over the network.

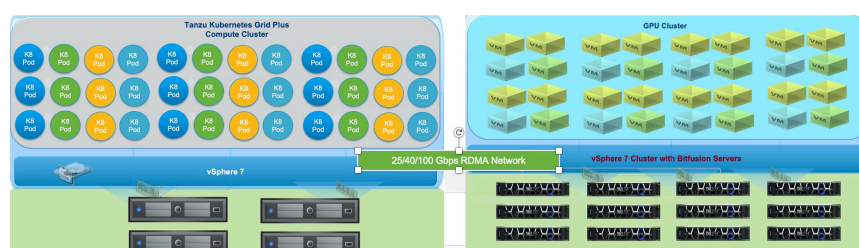


Figure 6: Logical architecture of the solution

Infrastructure components used in the solution are shown.

GPU Cluster	
Hosts	4 X Dell PowerEdge R740 servers
Storage	Infinidat F4300 Series Fibrechannel
GPUs	NVIDIA T4
Networking	Mellanox RoCE 100 Gb/s Ethernet

Table 1: GPU Cluster components used in the solution

CPU Cluster	
Hosts	4 X Dell PowerEdge R740 servers
Storage	Infinidat F4300 Series with Fibrechannel & NFS
Kubernetes	TKG Plus v16.6
Networking	25 Gb/s Ethernet

Table 2: CPU Cluster components used in the solution

5.1 VMware Tanzu Kubernetes Infrastructure:

Cloud native applications leverage VMware Tanzu to run Kubernetes natively on VMware. The initial setup of the scalable



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

solution leverages one master node and four worker nodes. The profile of the nodes are shown in figure.

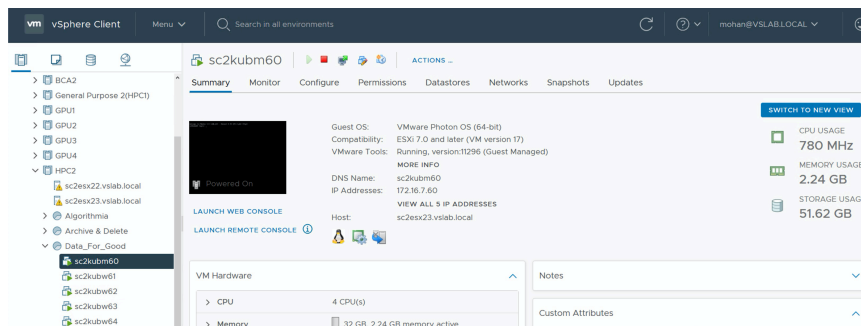


Figure 7: Virtual machines representing the Tanzu Kubernetes

The Kubernetes nodes and the pods used with Kubeflow are shown below. The Kubeflow namespace and its pods, which is deployed by vMPLP provides the capabilities of the machine learning platform used by the data scientists.

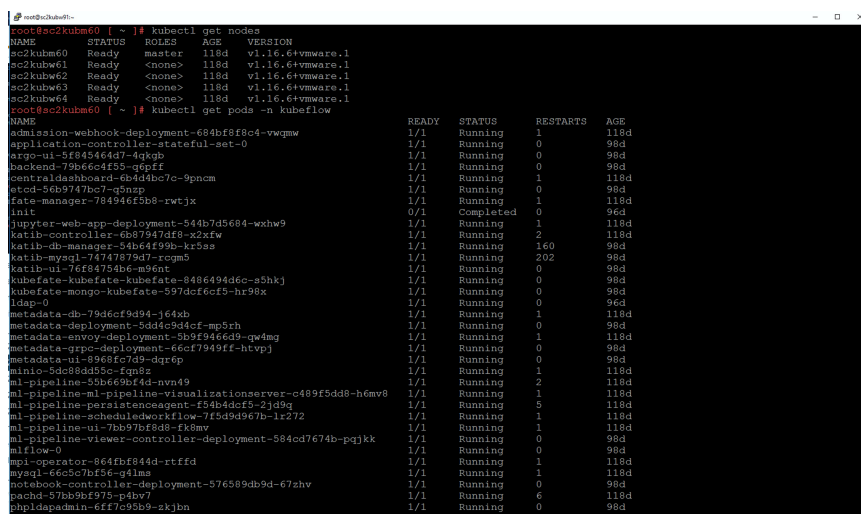


Figure 8: Kubernetes nodes and Kubeflow pods

5.2 GPU Cluster:

GPU resources are consolidated in a dedicated GPU cluster leveraging NVIDIA vGPU for local access and Bitfusion for remote access over the network. The GPU hosts have been setup with VMWare Bitfusion software for GPU virtualization



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

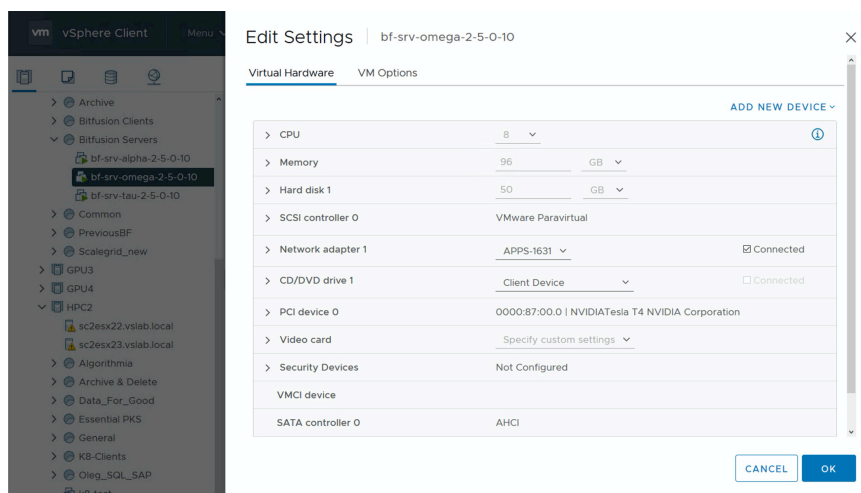


Figure 9: Virtual machine components of a Bitfusion GPU server

Bitfusion GPU resources can be monitored and managed via the vCenter plugin. Client and server components are seen in the dashboard.

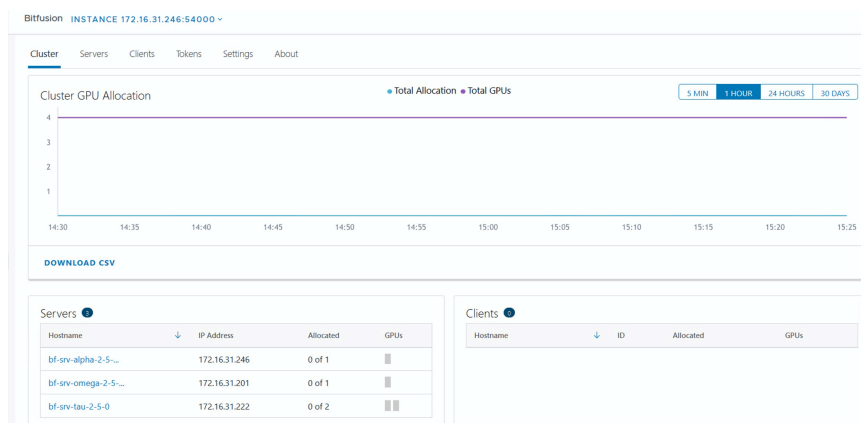


Figure 10: vCenter Bitfusion plugin for monitoring GPU resources

6 Solution in Action

The solution was successfully deployed in the VMware cloud Solutions Lab (VCSLAB) leveraging its state of the art hardware



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

and software components. The deployed solution provides a Jupyter hub interface for data scientists to bring their data and perform data modeling and analysis.

The primary components of the platform leveraging Kubeflow are shown below.

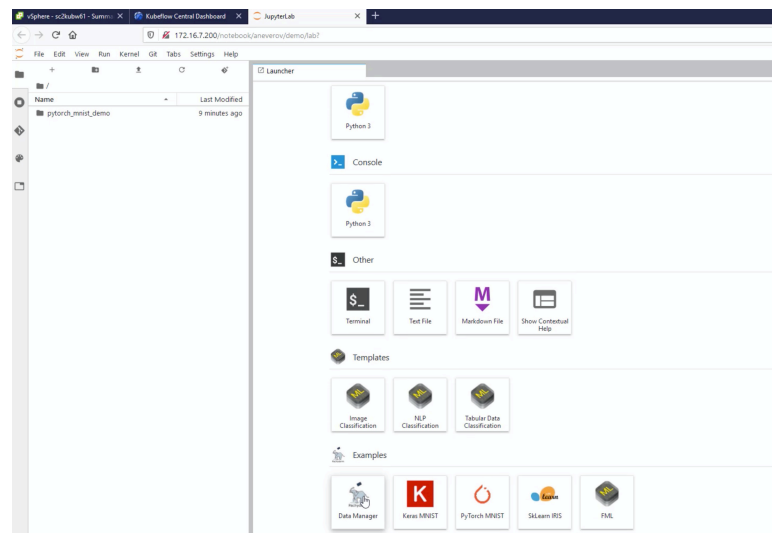


Figure 11: Components of the vMLP solution for data scientists

The data scientist has the capability to import their own data into the environment for their analysis.



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

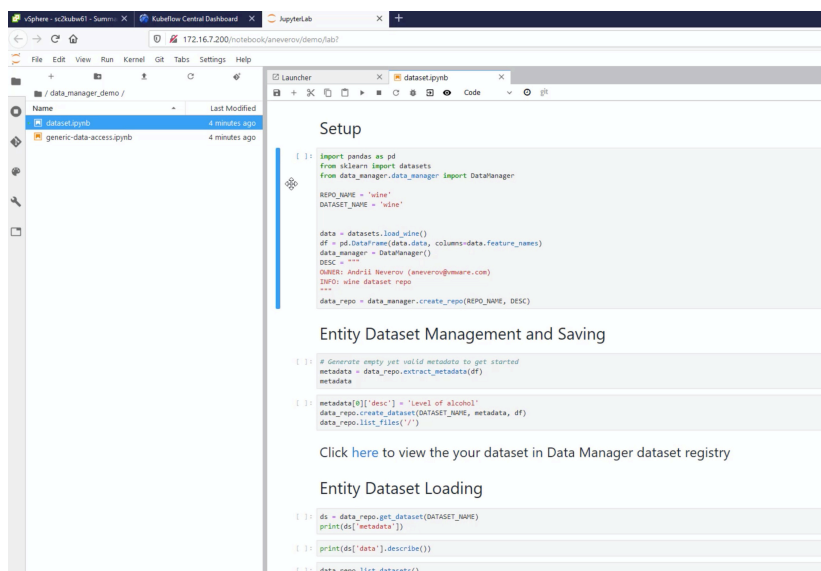


Figure 12: Importing data into the platform

After the data is imported, it can be modeled with machine learning frameworks like PyTorch as shown

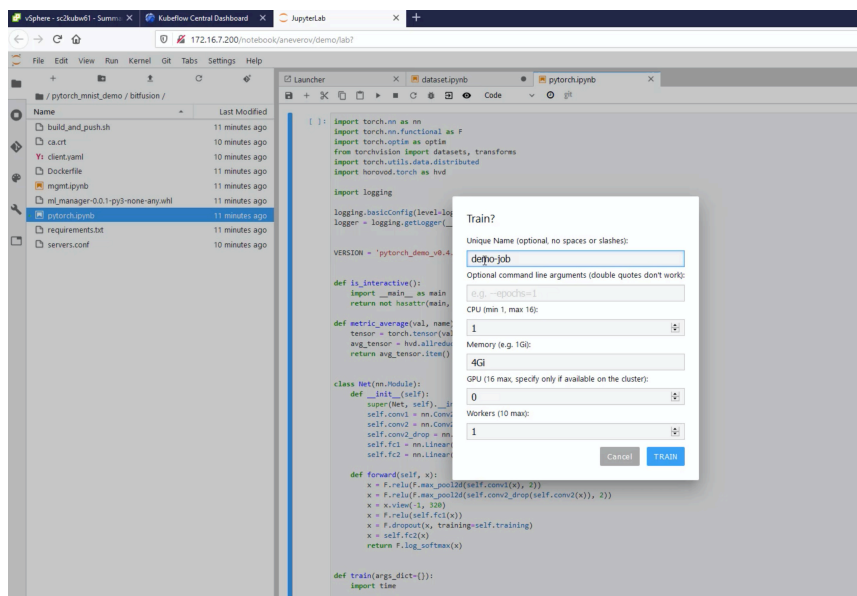


Figure 13: Modeling the data with PyTorch

The data model leverages GPUs in the backend through network access with Bitfusion.



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

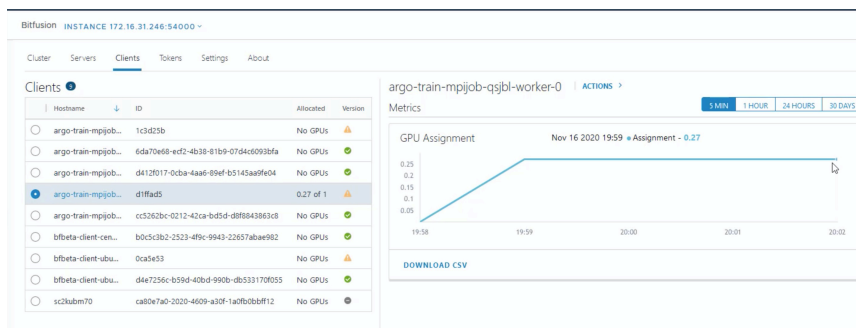


Figure 14: GPU utilization during model training

By iterating through multiple models, the data scientists eventually come up with the best fit. This model can then be packaged and deployed for inference in production.

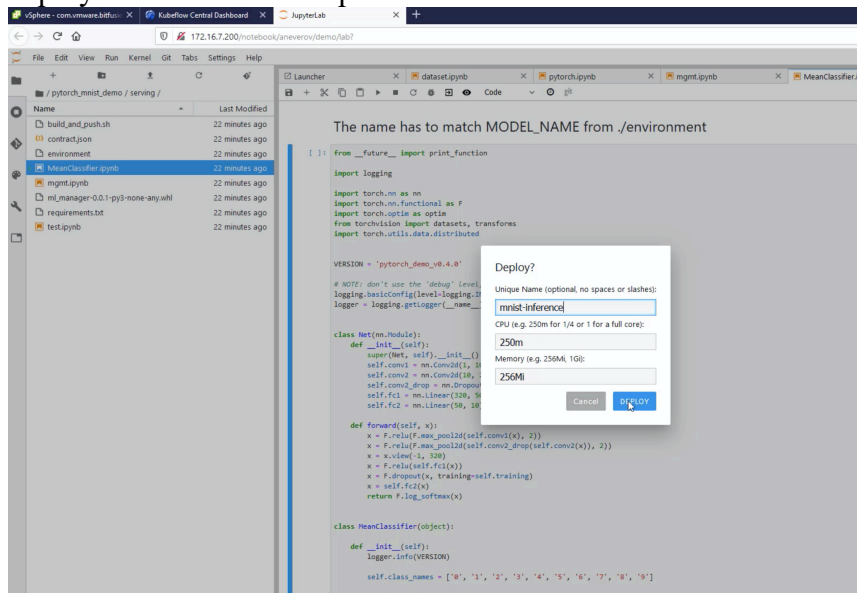


Figure 15: Deploying trained models for inference

7 Conclusion:

By building out this real world solution, we have effectively shown the following:

- VMware Cloud Foundation can be used as a force for good combines VMware Cloud
- Machine Learning platforms like VMLP can effectively run on VCF



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

- GPUs can be leveraged to speed up data science processing many fold
- VMware Tanzu based Kubernetes can be used to run ML/AI
- Data for Good Solution is a good proof point for running AI/ML on VCF

VMware Customers and VCPP partners can effectively leverage VCF to deploy a productive data science platform for their end users and customers.



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.