





## Accelerate performance on machine learning workloads with the Dell EMC PowerEdge C4140



This solution with NVIDIA Tesla V100 GPUs achieved more work than a Supermicro solution did and kept the GPUs cooler

More and more companies are using machine learning to solve business problems. To handle the massive amounts of compute power these workloads require, they are turning to servers that augment their CPUs with graphics processing units (GPUs) for additional compute power. We conducted hands-on testing of two such solutions, the Dell EMC™ PowerEdge™ C4140 and the Supermicro® SuperServer 1029GQ-TRT, both equipped with four NVIDIA® Tesla® V100 PCIe GPUs. Hardware that tolerates higher datacenter temperatures can allow companies to reduce their operational expenditures for cooling. For that reason, our tests explored not only how these two solutions handled a machine learning workload, but also how they fared under extremely warm datacenter conditions.

We learned that the Dell EMC solution delivered stronger machine learning performance than the Supermicro solution did and kept its GPUs cooler at both typical and unusually warm datacenter temperatures. The design of the Dell EMC PowerEdge C4140 places all four PCIe GPUs in the front of the server, with access to the coolest air. Also, the Dell EMC PowerEdge C4140 includes the integrated Dell Remote Access Controller (iDRAC) 9, which provides a great deal of granular information about the GPUs and real-time control over cooling fans. Together, these advantages make this platform an excellent choice for companies adopting machine learning or expanding their machine learning endeavors.



### Gets the job done

Delivered up to 26.3% greater GPU performance in a space-efficient 1U form factor



### Can take the heat

Kept GPUs up to 10.8 degrees cooler



### Gives you control

iDRAC 9 with Multi-Vector Cooling delivers granular GPU info and fan control in real time

## Machine learning: Coming soon to your business (if it hasn't already)

It's difficult to overestimate the impact of machine learning and other types of artificial intelligence on businesses of all kinds in upcoming years. In October 2017, research firm Gartner identified the top 10 strategic technology trends for 2018.<sup>1</sup> The top three trends were as follows:

1. AI foundation
2. Intelligent apps and analytics
3. Intelligent things

According to Gartner, "The ability to use AI to enhance decision making, reinvent business models and ecosystems, and remake the customer experience will drive the payoff for digital initiatives through 2025." A 2017 Gartner survey found that 59 percent of organizations were planning their AI strategies, while the remainder had already piloted or adopted AI solutions.<sup>2</sup>

One AI area growing particularly quickly is machine learning, where compute-intensive algorithms first train computers to perform certain tasks and then execute those tasks. Anyone with a smart phone that anticipates the next word of a sentence before they type it has experienced the power of machine learning, and the list of ways in which businesses can benefit from this set of technologies expands daily.



### Dell EMC works with NVIDIA to meet your needs

Dell EMC partners with NVIDIA to provide customers with multiple GPU and system configurations to take advantage of either PCIe Gen 3 or SXM2 form factors, depending on their workload. Both of our test servers used Tesla V100 PCIe GPUs.



## Stronger performance on machine-learning workloads

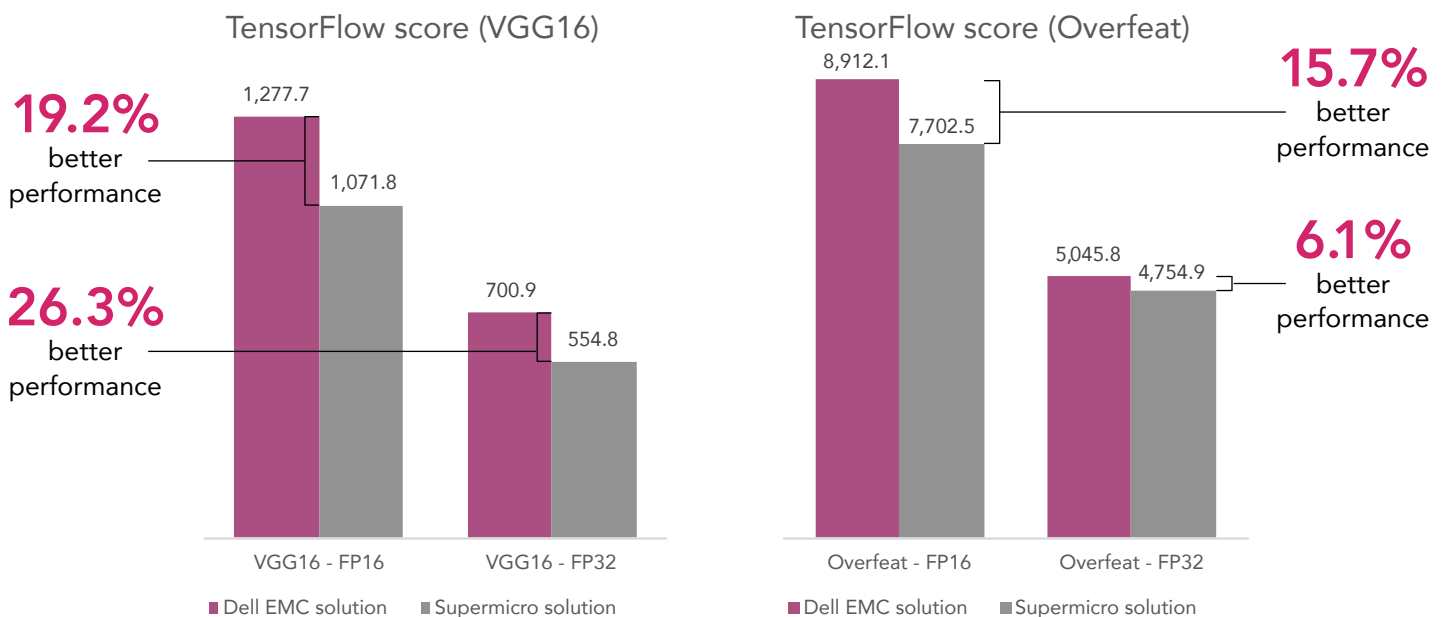
To measure the machine learning capabilities of the two solutions, we performed several image-classification workloads on each. Broadly speaking, image classification refers to computers learning to perform the very complex task of making sense of pictures using a type of algorithm called a convolutional neural network. Image classification has emerged as a valuable use case for machine learning.

Specifically, we used the `tf_cnn_benchmark`,<sup>3</sup> a tool that uses TensorFlow™, an open source framework for neural networks and other machine-learning applications. `tf_cnn_benchmark` includes different image-classification models you can run on either synthetic or real sample datasets containing images. To put it simply, neural network algorithms include two phases: training and inference. TensorFlow is used for both phases, but `tf_cnn_benchmark` tests only the training phase. Each test produces a result in terms of the number of images per second the system can process.

The models emphasize different aspects of machine learning and thus tax system components in different ways. We ran two models, VGG16 and Overfeat, on real sample datasets. Compared to the other models in `tf_cnn_benchmark`, both VGG16 and Overfeat require a relatively large amount of communication between the GPUs in a server. The Dell EMC PowerEdge C4140 is available in multiple configurations. One feature you can select is a PCIe switch, a high-speed interconnect that allows the GPUs in the server to communicate directly with each other without going through any CPU interface. The PowerEdge C4140 we tested had a PCIe switch, which gave it an advantage when performing these two inter-GPU communication-intensive workloads. The Supermicro SuperServer 1029GQ-TRT is not available with a PCIe switch.

We ran VGG16 and Overfeat two ways on each solution, with 16-bit floating point (FP) and with 32-bit FP. For all four tests, we used the ImageNet real data set<sup>4</sup> and used tuning settings that were optimal for each of the solutions under test. (For more details on our testing, see [Appendix B: How we tested.](#))

As the chart below shows, the Dell EMC solution outperformed the Supermicro solution on all four tests we conducted. Its greatest advantage was on VGG16 – FP32, where it processed 26.3 percent more images per second. This difference indicates the potential of the Dell EMC PowerEdge C4140 to perform either more machine learning work in a given period or the same amount of work in less time.



## If you can't stand the heat, get out of the datacenter

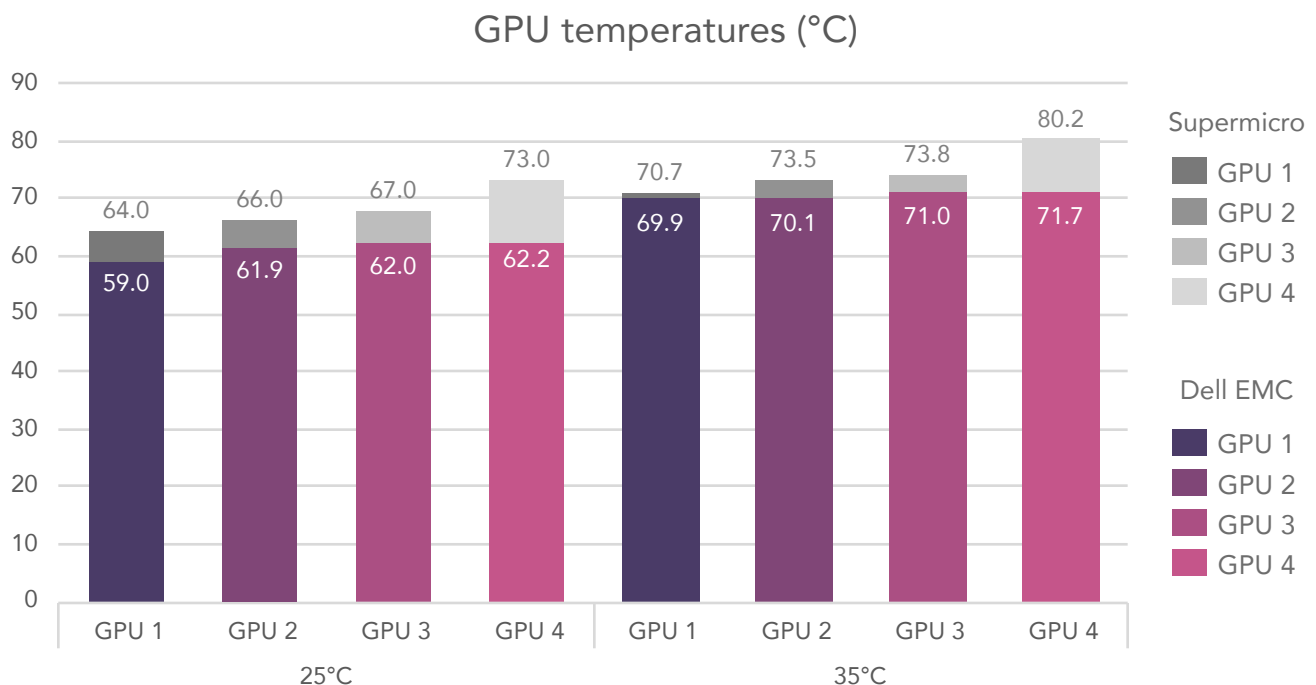
Companies spend dearly on datacenter cooling to avoid the potential performance degradation that can occur at excessive temperatures. By investing in hardware that operates with full functionality in warmer-than-normal datacenter conditions, they can set their thermostats higher and reduce this spending.

To learn how the two solutions performed at two different datacenter temperatures, 25°C and 35°C, we used the GPUBurn tool to ensure that the GPUs were under full load and the NVIDIA System Management Interface (SMI) program to read the temperature of the GPUs. When we raised the room temperature, we observed the following:

- Machine learning performance remained constant on both solutions.
- The power usage of both solutions increased as the servers expended more energy on cooling, but the total power usage at both 25°C and 35°C was lower on the Dell EMC solution than on the Supermicro solution. (See [Appendix C: Test results](#) for more details.)
- For both solutions, the average temperature of the four GPUs in the server increased, but the Dell EMC PowerEdge C4140 kept its GPUs cooler at both ambient temperatures than the Supermicro SuperServer 1029GQTRT did.

As the chart below shows, in a datacenter cooled to 25°C, the four GPUs in the Dell EMC solution averaged 61.3°C. When we raised the datacenter temperature to 35°C, the average temperature of the four GPUs increased to 70.7°C. In contrast, the GPUs in the Supermicro solution averaged 67.5°C at the lower ambient temperature and the average increased to 74.6°C under the warmer room conditions. Put another way, the NVIDIA GPUs in the Dell EMC solution were, on average, 6.2 degrees cooler than those in the Supermicro solution at an ambient temperature of 25°C and 3.9 degrees cooler at an ambient temperature of 35°C.

The Dell EMC solution's ability to keep its GPUs cooler, which we attribute to the cooling features we describe on the following page, would give a business the option of keeping its datacenter slightly warmer to save on cooling costs. Lower operating temperatures also have the potential to extend the functional life of hardware components.



## How the Dell EMC solution and the Supermicro solution differ when it comes to cooling

On the previous page, we presented our findings regarding the superior cooling of the Dell EMC solution over the Supermicro solution in our testing. We attribute that advantage to two primary differences between the solutions: their physical design and their management features.

### GPUs stay cooler in front

In the Supermicro, three of the four GPUs are in the front of the server, where airflow is greatest, and the fourth GPU is in the back of the server. In contrast, the design of the Dell EMC PowerEdge C4140 accommodates all four GPUs in the front. Our results reflect this difference. Under both ambient conditions, the temperature of the four GPUs in the Dell EMC solution varied by less than 4 degrees Celsius. In the Supermicro, three of the four GPUs—the ones located in the front—also varied by less than 4 degrees. However, compared to the average temperature of the three GPUs in the front, the fourth GPU in the back of the system was 4.6 degrees warmer at the lower ambient temperature and 7.5 degrees warmer at the higher ambient temperature.

### Management matters

The Dell EMC PowerEdge C4140 offers several management features that can help administrators control the temperature of the server and its GPUs. Its fully featured Baseboard Management Controller (BMC) interface includes an airflow measurement sensor readout and can change fan speed based on current airflow. This readout makes it easier to optimize datacenter cooling and airflow requirements to keep costs down. It also offers a Multi-Vector Cooling feature in iDRAC 9 that allows the user to customize fan speed based on desired linear feet per minute (LFM) so that a business can optimize its datacenter power and cooling needs. Multi-Vector Cooling provides three options:

- Maximum Performance (Performance Optimized): for reduced memory and CPU throttling and higher fan speed.
- Minimum Power (Performance per Watt Optimized): for decreased power consumption.
- Sound Cap: for reduced acoustical output at the expense of some performance.

Users can fine-tune the baseline Thermal Profiles, or fan control profiles, with the following settings:

- Fan Speed Offset (Low/Med/High/Max)
- Maximum Exhaust Temperature Limit
- Minimum Fan Speed

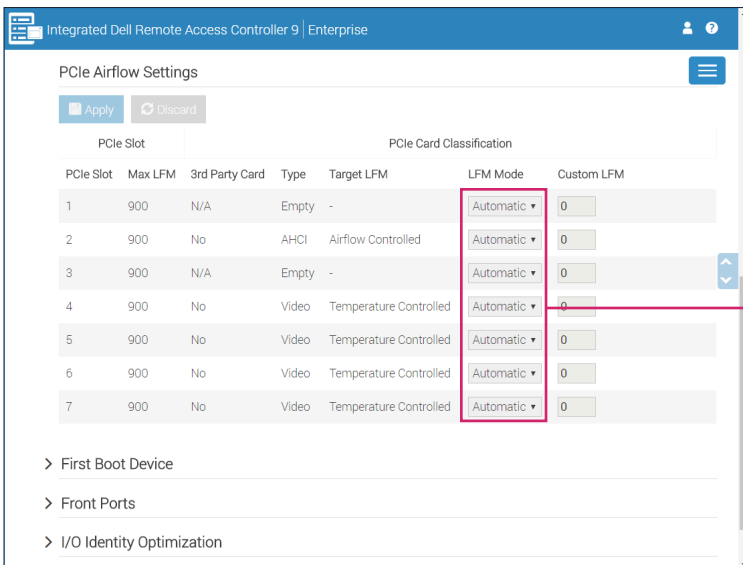
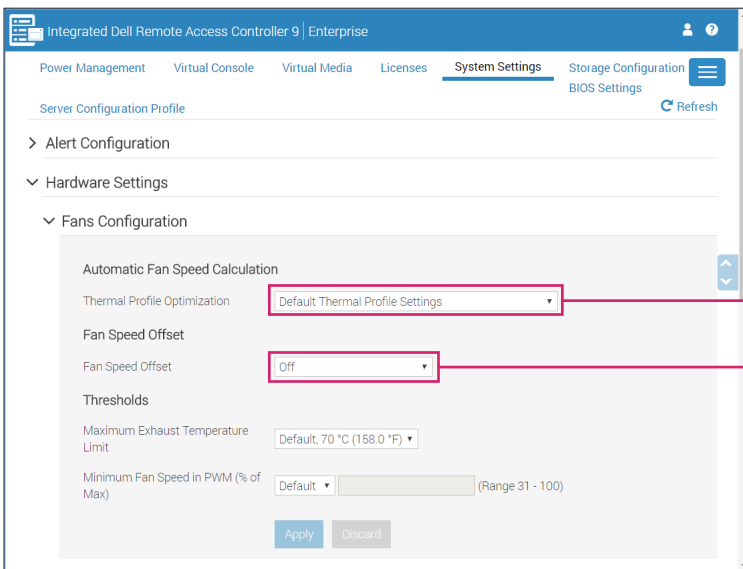
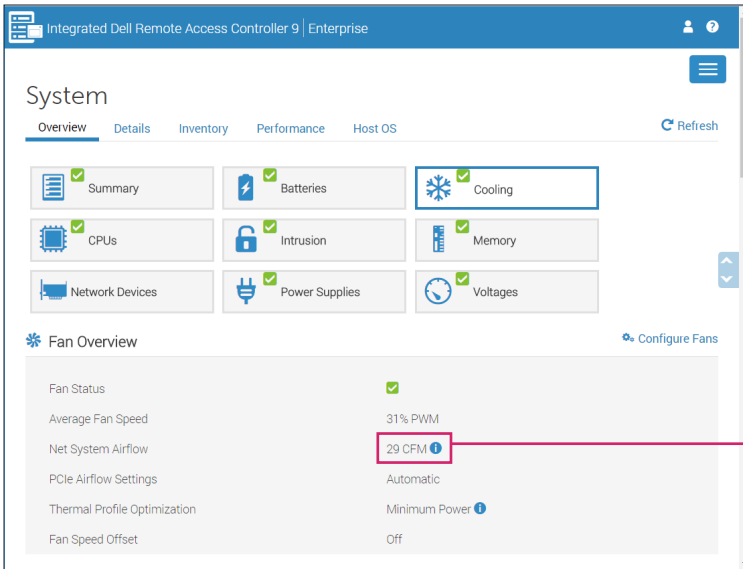
While the settings listed above apply to the server as a whole, users can set individual PCIe slots to custom airflow settings on a slot-by-slot basis. Each PCIe slot can automatically adjust airflow based on information on the card, or a custom LFM value. Users can also disable PCIe slot airflow entirely. Doing so would save power for a card with built-in cooling capabilities (e.g., an onboard fan).

The Supermicro solution lacks the control and insight that come with iDRAC 9 Multi-Vector Cooling. It has a relatively rudimentary BMC interface that allows for fan speed adjustment but only with basic low-medium-high settings. It also lacks the iDRAC 9 interface's fine-tuning abilities, such as being able to adjust fan speed to meet a specific airflow requirement. The Supermicro BMC also lacks an airflow measurement sensor readout and the ability to adjust thermal profiles directly from the BMC.



## Knowledge is power: How iDRAC 9 management tools give you information and control

The Multi-Vector Cooling feature included in iDRAC 9 gives you more data and granular control over the cooling of your PowerEdge C4140 server, as you can see in the screenshots we took while testing. The first screen shows the real-time airflow sensor readout, which lets you understand the current airflow conditions from the comfort of a management interface. The second screen shows the fan speed control options, with baseline Thermal Profile settings for performance, power, or acoustical requirements, and additional fine-tuning options for exhaust temperature and minimum fan speed limits. The third screen shows the slot-by-slot PCIe airflow optimizations, which you can use to configure airflow based on an individual PCIe slot's hardware needs.



Real-time airflow sensor readout

Can be Default / Maximum Performance / Minimum Power / Sound Cap

Can be Off / Low / Med / High / Max

Can be Automatic / Custom LFM / Disabled



## Conclusion

Whether your company has already begun to use machine learning or is still gathering information for a future machine learning initiative, you'll want to understand the factors that differentiate the various solutions on the market. The fact that a server accommodates four powerful NVIDIA Tesla V100 graphics processing units to augment the processing power of its CPUs is an important starting point. But what levels of performance can it deliver? What features does it have to keep GPUs cool so you can set your datacenter thermostat a bit higher to help reduce your power bill? And how can its management tools add value to your business?

In our hands-on tests, the Dell EMC PowerEdge C4140 with Multi-Vector Cooling outperformed the Supermicro SuperServer 1029GQ-TRT by up to 26.3 percent on a machine learning workload. It also kept its four NVIDIA Tesla GPUs up to 10.8 degrees Celsius cooler. The server's iDRAC 9 interface provided more system information and a greater degree of control over its cooling features than the Supermicro management did. These factors make the Dell EMC PowerEdge C4140 a choice worth considering as your company rides the machine learning wave.

- 
- 1 Kasey Panetta, "Gartner Top 10 Strategic Technology Trends for 2018," accessed March 21, 2018, <https://www.gartner.com/smarterwithgartner/gartner-top-10-strategic-technology-trends-for-2018/>
  - 2 Kasey Panetta, "Gartner Top 10 Strategic Technology Trends for 2018," accessed March 21, 2018, <https://www.gartner.com/smarterwithgartner/gartner-top-10-strategic-technology-trends-for-2018/>
  - 3 Learn more at [https://github.com/tensorflow/benchmarks/tree/master/scripts/tf\\_cnn\\_benchmarks](https://github.com/tensorflow/benchmarks/tree/master/scripts/tf_cnn_benchmarks)
  - 4 Learn more at <http://www.image-net.org/>



On April 20, 2018, we finalized the hardware and software configurations we tested. Updates for current and recently released hardware and software appear often, so unavoidably these configurations may not represent the latest versions available when this report appears. For older systems, we chose configurations representative of typical purchases of those systems. We concluded hands-on testing on June 15, 2018.

## Appendix A: System configuration information

Server configuration information	Dell PowerEdge C4140	Supermicro SuperServer SYS-1029GQ-TRT
BIOS name and version	Dell BIOS 1.1.7	American Megatrends 2.0b
Non-default BIOS settings	N/A	Onboard RAID enabled
Operating system name and version/ build number	CentOS 7.5.1804 kernel-3.10.0-862.3.2	CentOS 7.5.1804 kernel-3.10.0-862.3.2
Date of last OS updates/patches applied	06/13/2018	06/14/2018
Power management policy	Performance	Max Performance
PCIe switch	Yes	No
Processor		
Number of processors	2	2
Vendor and model	Intel® Xeon® Gold 6148	Intel Xeon Gold 6148
Core count (per processor)	20	20
Core frequency (GHz)	2.4	2.4
Stepping	4	4
Memory module(s)		
Total memory in system (GB)	192	192
Number of memory modules	12	12
Vendor and model	Hynix HMA82GR7AFR8N-VK	Hynix HMA82GR7AFR8N-VK
Size (GB)	16	16
Type	PC4-2666	PC4-2666
Speed (MHz)	2,666	2,666
Speed running in the server (MHz)	2,666	2,666
Graphics processing units		
Number of drives	4	4
Drive vendor and model	NVIDIA Tesla V100-PCIE-16GB	NVIDIA Tesla V100-PCIE-16GB
Driver version	390.30	390.30





Server configuration information	Dell PowerEdge C4140	Supermicro SuperServer SYS-1029GQ-TRT
Local storage - Type 1		
Number of drives	2	2
Drive vendor and model	Intel SSDSCKJB240G7R	Intel SSDSA2BZ10
Drive size (GB)	240	100
Drive information (speed, interface, type)	SATA SSD	SATA SSD
Local storage - Type 2		
Number of drives	1	1
Drive vendor and model	Intel SSDPEDMD020T401	Intel SSDPEDMD020T401
Drive size (TB)	2	2
Drive information (speed, interface, type)	PCIe SSD	PCIe SSD
Network adapter		
Vendor and model	Intel X540/I350	Intel X540-AT2
Number and type of ports	4 RJ45	2 RJ45
Driver version	4.4.0-k-rh7.4	4.4.0-k-rh7.4
Cooling fans		
Vendor and model	Delta Electronics® GFM0412SS-D	SUNON® PF40561BX-Q40U-S9H
Number of cooling fans	8	9
Power supplies		
Vendor and model	Dell 0W1R7V	Supermicro PWS-2K05A-1R
Number of power supplies	2	2
Wattage of each (W)	2,000	2,000



# Appendix B: How we tested

## Installing and configuring the operating system

After setting the BIOS settings to Performance on the Dell EMC server and Max Performance on the Supermicro server, we installed CentOS 7 onto each server. We updated the operating system and installed prerequisites for the NVIDIA Docker container. We performed the following steps on each server:

### Changing the BIOS settings

On the Dell EMC server, set the power profile to Performance.

On the Supermicro server, set the power management to Max Performance as follows:

Advanced→CPU Configuration→Advanced Power Management Configuration→

- Power Technology: Custom (Default: Energy Efficient)
- Power Performance Tuning: BIOS Controls EPB (Default: OS Controls EPB)
- Energy Performance BIAS Setting: Max Performance (Default: Balanced Performance)

### Installing CentOS 7

1. Boot the server to the CentOS 7 installation media.
2. Select Install CentOS 7, and press Enter.
3. Select English (United States) as the installation language, and click Continue.
4. At the Installation Summary screen, select the following options:
  - a. Set the Date and Time settings to the local time zone.
  - b. Set the Software Selection to Minimal Install.
  - c. Set the Installation Destination as Local Disk, and Automatically configure partitioning.
  - d. Set the Network and Hostname to DHCP, and turn the Ethernet device on.
5. Click Begin Installation.
6. Enter a root password, and click Done twice.

### Updating and configuring the operating system

1. Log into the operating system as the root user.
2. Disable SELinux and the firewall by running the following commands:

```
sed -i 's/SELINUX=.*SELINUX=disabled/' /etc/selinux/config
systemctl disable firewalld
```
3. Update the operating system and reboot by running the following commands:

```
yum -y update
reboot
```
4. Install the following packages:

```
yum -y install xfsprogs tuned wget vim nfs-utils openssh-clients man zip unzip numactl sysstat
pciutils
```
5. Configure Tuned:

```
tuned-adm profile throughput-performance
```
6. Install prerequisites and prepare system for the NVIDIA drivers:

```
echo "blacklist nouveau" > /etc/modprobe.d/blacklist-nouveau.conf
echo "options nouveau modeset=0" >> /etc/modprobe.d/blacklist-nouveau.conf
dracut -force
reboot
```
7. Install the NVIDIA drivers:

```
wget http://us.download.nvidia.com/tesla/390.30/NVIDIA-Linux-x86_64-390.30.run
sh NVIDIA-Linux-x86_64-390.30.run
```
8. Enable NVIDIA Persistence Mode on boot:

```
echo "/usr/bin/nvidia-persistenced --verbose" >> /etc/rc.d/rc.local
chmod +x /etc/rc.d/rc.local
reboot
```



## Installing the NVIDIA Docker container

We used the NVIDIA Docker container prebuilt with TensorFlow to run the benchmark. Perform the following steps on each server:

1. Install Docker CE with the following commands:

```
yum install -y yum-utils device-mapper-persistent-data lvm2
yum-config-manager --add-repo https://download.docker.com/linux/centos/docker-ce.repo
yum install -y docker-ce
systemctl start docker
systemctl enable docker
```
2. Install NVIDIA Docker:

```
curl -s -L https://nvidia.github.io/nvidia-docker/centos7/x86_64/nvidia-docker.repo | tee /etc/yum.
  repos.d/nvidia-docker.repo
yum install -y nvidia-docker2
systemctl restart docker
```
3. Create and run the TensorFlow image:

```
nvidia-docker create --name=tf_test --shm-size=1g --ulimit memlock=-1 --ulimit stack=67108864 --
  privileged --mount type=bind,src=tf_test,dst=tf_test -i tensorflow/tensorflow:1.7.1-devel-gpu
nvidia-docker start -i tf_test
```
4. Rebuild Tensorflow to support all CPU extensions:

```
export TF_CUDA_COMPUTE_CAPABILITIES=7.0
cd /tensorflow
ln -s /usr/local/cuda/lib64/stubs/libcuda.so /usr/local/cuda/lib64/stubs/libcuda.so.1 && \
LD_LIBRARY_PATH=/usr/local/cuda/lib64/stubs:${LD_LIBRARY_PATH} \
tensorflow/tools/ci_build/builds/configured GPU \
bazel build -c opt --config=cuda \
--cxxopt="-D_GLIBCXX_USE_CXX11_ABI=0" \
--copt=-march="native" \
tensorflow/tools/pip_package/build_pip_package && \
/usr/local/cuda/lib64/stubs/libcuda.so.1 && \
bazel-bin/tensorflow/tools/pip_package/build_pip_package /tmp/pip && \
pip --no-cache-dir install --upgrade /tmp/pip/tensorflow-*.whl && \
rm -rf /tmp/pip && \
rm -rf /root/.cache
```

## Installing and running the benchmark

We pulled the benchmark `tf_cnn_benchmark` from TensorFlow's Github repo. We made some changes to the run length and sample size to ensure steady-state GPU temperature. To make these changes, run the following commands inside the `nvidia-docker` container on each server:

1. Pull `tf_cnn_benchmark`:

```
cd /tensorflow
git clone https://github.com/tensorflow/benchmarks.git
cd benchmarks/
git checkout bab8a61aaca3d2b94072ae2b87f0aafef1797b165
cd benchmarks/scripts/tf_cnn_benchmarks/
```
2. Run the benchmark with the following commands.

All tests:

```
python tf_cnn_benchmark.py --data_name=imagenet --data_dir=/data/imagenet --
  num_batches=100 --num_gpus=4 --data_format=NCHW --optimizer=sgd --distortions=False
```

Switching between FP16 and FP32:  
FP16:  
FP32: `--use_fp16=False`

Server- and model-specific tuning:

Dell EMC solution:

```
--model=vgg16 --batch_size=128 --use_fp16=True --local_parameter_device=gpu --
  variable_update=replicated --all_reduce_spec=nccl
--model=vgg16 --batch_size=64 --use_fp16=False --local_parameter_device=gpu --
```



```
variable_update=replicated --all_reduce_spec=nccl
--model=overfeat --batch_size=2048 --use_fp16=True --local_parameter_device=gpu --
variable_update=replicated --all_reduce_spec=nccl
--model=overfeat --batch_size=1024 --use_fp16=False --local_parameter_device=gpu --
variable_update=replicated --all_reduce_spec=nccl
```

Supermicro solution:

```
--model=vgg16 --batch_size=128 --use_fp16=True --local_parameter_device=cpu --
variable_update=parameter_server --all_reduce_spec=
--model=vgg16 --batch_size=64 --use_fp16=False --local_parameter_device=gpu --
variable_update=replicated --all_reduce_spec=
--model=overfeat --batch_size=2048 --use_fp16=True --local_parameter_device=gpu --
variable_update=parameter_server --all_reduce_spec=
--model=overfeat --batch_size=1024 --use_fp16=
```

## Preparing and running GPU Burn

1. Run the following commands to prepare the GPU Burn container:

```
echo `yum update -y && yum install -y make && curl http://wili.cc/blog/entries/gpu-burn/
gpu_burn-0.9.tar.gz | tar -xz && sed -i `s/compute_30/compute_70/` Makefile && sed -i -e `s/float.
*nextReport *= *[0-9].*/float nextReport = 1.0f;/` -e `s/nextReport *= *elapsed *+ *[0-9].*/
nextReport
= elapsed + 1.0f;/` gpu_burn-drv.cpp && make && exit' | nvidia-docker run --name=gpu_burn -i
nvidia/cuda:9.1-devel-centos7
nvidia-docker start gpu_burn
```

2. Use the following command to run gpu\_burn:

```
nvidia-docker exec gpu_burn /gpu_burn -d 1800
```



## Appendix C: Test results

In the tables below, we present the results of our performance testing using the tf\_cnn\_benchmark tool and our thermal testing, where we measured the temperature of the GPUs while the system was under load.

tf_cnn_benchmark	Dell EMC solution	Supermicro solution	Difference (percentage)
VGG16 - FP16	1277.78	1071.89	19%
VGG16 - FP32	700.97	554.89	26%
Overfeat - FP16	8912.13	7702.55	16%
Overfeat - FP32	5045.88	4754.92	6%

Temperature testing	Dell EMC solution	Supermicro solution	Difference (degrees)
Ambient temperature 25°C			
GPU1	59.0	64.0	5.0
GPU2	61.9	66.0	4.1
GPU3	62.0	67.0	5.0
GPU4	62.2	73.0	10.8
Average	61.3	67.5	6.2
Ambient temperature 35°C			
GPU1	69.9	70.7	0.8
GPU2	70.1	73.5	3.4
GPU3	71.0	73.8	2.8
GPU4	71.7	80.2	8.5
Average	70.7	74.6	3.9

Power testing (W)	Dell EMC solution	Supermicro solution	Difference (percentage)
Power usage at 25°C	1378.4	1406.1	1.97%
Power usage at 35°C	1467.3	1517.3	3.30%



This project was commissioned by Dell EMC.



Facts matter.®

Principled Technologies is a registered trademark of Principled Technologies, Inc.  
All other product names are the trademarks of their respective owners.

**DISCLAIMER OF WARRANTIES; LIMITATION OF LIABILITY:**

Principled Technologies, Inc. has made reasonable efforts to ensure the accuracy and validity of its testing, however, Principled Technologies, Inc. specifically disclaims any warranty, expressed or implied, relating to the test results and analysis, their accuracy, completeness or quality, including any implied warranty of fitness for any particular purpose. All persons or entities relying on the results of any testing do so at their own risk, and agree that Principled Technologies, Inc., its employees and its subcontractors shall have no liability whatsoever from any claim of loss or damage on account of any alleged error or defect in any testing procedure or result.

In no event shall Principled Technologies, Inc. be liable for indirect, special, incidental, or consequential damages in connection with its testing, even if advised of the possibility of such damages. In no event shall Principled Technologies, Inc.'s liability, including for direct damages, exceed the amounts paid in connection with Principled Technologies, Inc.'s testing. Customer's sole and exclusive remedies are as set forth herein.