

COMPUTE INTENSIVE PERFORMANCE EFFICIENCY COMPARISON: HP MOONSHOT WITH AMD APUS VS. AN INTEL PROCESSOR-BASED SERVER

HP Moonshot with ProLiant m700 cartridges powered by **AMD Opteron™ X2150 APUs**



12.6X greater job throughput
for compute-intensive workloads

up to **12.7%** more power efficient
measured in kWh/operation

versus an Intel® Xeon® Processor E5-2660 v2-powered server



Increased use of highly-parallel architectures for compute intensive workloads, such as render farms, has led to the development of a new class of products that unify graphics processing and general computation, such as AMD's accelerated processing unit (APU) offerings. One of the main benefits provided by AMD's integration of graphics and computing technologies is the power efficiencies achieved by these products. Sharing computational, graphics, and chip data path resources help to reduce power consumption per compute operation and provide improved performance efficiencies. Another potential benefit of the APU is reducing total cost of ownership (TCO) for businesses running workloads where data processing, graphics, and visualization all play an important role, such as graphics-based applications, hosted desktops, and image processing and rendering. Finally, an important factor to consider along with APU benefits is the form factor of the physical servers that you choose for your compute-intensive workload, because rack space savings in the data center can lead to lower operational expenses in cooling and other infrastructure.

In the Principled Technologies labs, we performed a compute intensive workload, 3D rendering tasks, on two platforms: an AMD-based HP Moonshot 1500 chassis filled with HP ProLiant m700 server cartridges and an Intel Xeon processor E5-2660 v2-based server. The HP Moonshot solution provided over 12 times the job throughput of a single Intel server, meaning it would take more than 12 Intel servers to



accomplish the same work in the same time, and it was also more power efficient per workload operation than the Intel solution. Finally, it accomplished the work in a 4.3U rack space, as opposed to the 12U that 12 Intel servers would have consumed.

COMPUTE-INTENSIVE ENERGY EFFICIENCY

One use case for APUs that takes advantage of AMD's technology advancements are highly parallel compute intensive workloads, such as high-quality graphics rendering. In the case of graphics rendering, as rendering needs become more compute intensive, it is a challenge to render frames in a reasonable amount of time while also using energy as efficiently as possible across the data center. Building out a compute farm is an expensive solution, and not only from the hardware standpoint. Space, cooling, and power capabilities limit many data centers, making it challenging to simply throw more machines at the problem.

Moving computation to Internet-based cloud computing providers has downsides—increased cost, significant bandwidth requirements, and potential concerns regarding security. Performing this computation on traditionally non-dense form factors gets the job done, but at higher overhead and power costs. A solution to this problem is to use massive parallelization via energy-efficient low power APUs from AMD in ultra-high density environments. This type of solution allows for over a thousand nodes per rack in some configurations. One of the first solutions based on this model is the AMD-based HP Moonshot 1500 chassis with the HP ProLiant m700 server cartridge, which allows up to 1,800 AMD Opteron™ X2150 APUs in a full-rack configuration.

We tested this HP Moonshot system filled with 45 HP ProLiant m700 server cartridges (180 total APUs), to understand how these new APU-based computing systems compare to traditional server architectures in terms of processing efficiency and power consumption. See [Appendix B](#) for information about our test systems and [Appendix C](#) for how we set up and ran the tests.

WHAT WE FOUND

About the results

We measured rendering rates or job throughput, along with energy consumption, for the HP Moonshot system with 45 ProLiant m700 server cartridges and for the Intel server.

- In our tests, the HP Moonshot with ProLiant m700 cartridges delivered **12.6 times greater job throughput** for the 3D rendering workloads than with a single Intel system. Although the Intel system could render more quickly (3.4 to 2.4 times faster than one ProLiant m700 node for the cases we considered), the 180 AMD-powered nodes sped the job up.

- Depending on the way we ran the workload on the Intel system, the HP Moonshot with ProLiant m700 cartridges consumed from **10.0 to 12.7 percent less energy**, measured in kWh, than one Intel system while performing the same amount of work.

Throughput, or total rate, is the number of rendering operations per second for the entire system. A higher throughput is better, as a higher rate means the system can perform more work.

Energy consumption depends on the average power over the run divided by the job throughput, so we report energy use in kilowatt-hours used by the system per rendering operation.

We investigated system performance for a variety of system loads: we varied the number of identical instances of the 3D rendering program and the number of CPU threads assigned to each program. For the Moonshot, we found the best performance for four CPU threads and all GPU threads.

| HP Moonshot | One HP ProLiant m700 node (average) | | | | System (180 nodes) | | |
|------------------------|-------------------------------------|---------|---------------|--|--------------------------------------|---------------|-------------------------------|
| Subscription (threads) | Number of instances | Threads | Total threads | Throughput - Rate per instance (OPs/s) | Throughput - Rate per system (OPs/s) | Power (Watts) | System Energy (kWh/system/OP) |
| 100% | 1 | 4 | 4 | 29.8 | 5,368.8 | 3,264.3 | 1.69x10 ⁻⁷ |
| Intel server | One Intel server (average) | | | | System (1 Intel server) | | |
| Subscription (threads) | Number of instances | Threads | Total threads | Throughput - Rate per instance (OPs/s) | Throughput - Rate per system (OPs/s) | Power (Watt) | System Energy (kWh/system/OP) |
| 100% | 4 | 10 | 40 | 101.3 | 405.3 | 282.6 | 1.94x10 ⁻⁷ |
| 120% | 6 | 8 | 48 | 70.8 | 425.0 | 287.2 | 1.88x10 ⁻⁷ |

Figure 1: Performance and energy consumption for the two platforms. Greater throughput is better and lower energy consumption is better.

CONCLUSION

AMD’s accelerated processing units can be an enormous boon to those who perform compute intensive processing workloads, such as the 3D rendering workload we tested. In the Principled Technologies labs, an AMD-based HP Moonshot 1500 chassis with the ProLiant M700 server cartridge outperformed an Intel Xeon processor E5-2660 V2-based server —delivering 12.6 times the rendering performance of a single Intel server. It achieved this performance advantage while utilizing 10 percent less power than the more traditional server solution, and used just 4.3U of rack space instead of the 12U that 12 Intel servers would have used.

APPENDIX A – ABOUT THE COMPONENTS

About the HP Moonshot

According to HP, the Moonshot System with ProLiant m700 Server Cartridges “offers up to 44% lower TCO, while dramatically improving security and compliance by centralizing desktops, data, and applications in the data center. With four AMD Opteron X2150 APUs per cartridge, the ProLiant m700 Server Cartridge delivers up to 720 processor cores/chassis along with fully-integrated graphics processing to enhance productivity from any device. Up to 45 server cartridges fit in one converged Moonshot System for 1,800 servers/rack, so you spend less on racks, switches and cables.”

Learn more at h17007.www1.hp.com/us/en/enterprise/servers/products/moonshot/

About the workload

LuxRender is a physically based rendering engine that simulates the flow of light according to physical equations, which lets it produce realistic images of photographic quality. In our testing, we used LuxRender with the an exporter plug-in for Blender 2.6x, LuxBlend 2.5. According to LuxRender, “LuxBlend 2.5 exposes virtually all LuxRender features and offers tight Blender integration via the binary pylux module.”

Learn more at www.luxrender.net/en_GB/index

APPENDIX B – SYSTEM CONFIGURATION INFORMATION

Figure 2 provides detailed configuration information for the test systems.

| System | HP ProLiant m700 Server Cartridge | Intel white box Supermicro® 6017R-WRF |
|---------------------------------------|---|--|
| General | | |
| Number of processor packages | 4 | 2 |
| Number of cores per processor | 4 | 10 |
| Number of hardware threads per core | 1 | 2 |
| Number of GPU cores per processor | 128 | N/A |
| Type of GPU cores | AMD Radeon 8000 | N/A |
| CPU | | |
| Vendor | AMD | Intel |
| Name | Opteron APU | Xeon |
| Model number | X2150 | E5-2660 v2 |
| Stepping | 1 | 04 |
| Core frequency (GHz) | 1.5 | LGA2011 |
| Bus frequency (MHz) | 800 | 2.20 |
| L1 cache | 192kB | 4000 |
| L2 cache | 4096kB | 640kB |
| L3 cache | N/A | 2.5MB |
| Chassis | | |
| Vendor and model number | HP Moonshot System | Supermicro 6017R-WRF |
| Motherboard model number | 1500 | X9DRW-iF |
| BIOS name and version | HP A34 | Intel C602 |
| BIOS settings | Preset to Balanced Power and Performance under OS Control | American Megatrends 3.0b |
| Memory module(s) | | |
| Total RAM in system (GB) | 32 | 128 |
| Vendor and model number | SK Hynix® HMT41GA7AFR8A-PB | Kingston® KVR16LR11D4/16KF |
| Type | PC3-12800 | PC3L-12800R |
| Speed (MHz) | 1,600 | 1,600 |
| Speed running in the system (MHz) | 1,600 | 1,333 |
| Timing/Latency (tCL-tRCD-tRP-tRASmin) | 11-11-11 | 11-11-11 |
| Size (GB) | 8 | 16 |
| Number of RAM module(s) | 4 | 8 |
| Chip organization | Double-sided | Double-sided |
| Rank | 2 | 2 |

| System | HP ProLiant m700 Server Cartridge | Intel white box Supermicro® 6017R-WRF |
|----------------------------|-----------------------------------|--|
| Operating system | | |
| Name | CentOS 6.5 x86_64 | CentOS 6.5 x86_64 |
| File system | ext4 | ext4 |
| Kernel | 2.6.32-431.11.2.el6.x86_64 | 2.6.32-431.11.2.el6.x86_64 |
| Language | English | English |
| Disk | | |
| Vendor and model number | ATA SanDisk SSD i110 | Seagate ST1000NM0033-9ZM173 |
| Number of disks in system | 4 | 2 |
| Size (GB) | 32 | 1,000 |
| Type | SATA, UDMA/133 | SATA 6 Gbs |
| Driver (Module) | lsg | N/A |
| Driver Version | 3.5.34 | N/A |
| Buffer size (MB) | N/A | 128 |
| RPM | N/A | 72,000 |
| Ethernet | | |
| Vendor and model number | Broadcom® NetXreme® BCM5720 | Intel Ethernet Server Adapter I350 Gigabit |
| Type | integrated | Integrated |
| Driver (Module) | tg3 | igb |
| Driver Version | 3.132 | 5.0.5-k |
| Power supplies | | |
| Total number | 3 (Moonshot chassis) | 2 |
| Vendor and model number | HP DPS-1200SB A | Supermicro PWS-704P-1R |
| Wattage of each (W) | 1200 | 700 |
| Cooling fans | | |
| Total number | 5 (Moonshot chassis) | 5 |
| Vendor and model number | Delta PFR0812XHE | Nidec® R40W12BS5AC-65 |
| Dimensions (h x w) of each | 8cmx8cmx3.8cm | 4cmx4cmx5.6cm |
| Volts | 12 | 12 |
| Amps | 4.9 | 0.84 |
| Disk controller | | |
| Vendor and model | N/A | IntelC600 Controller |
| Controller Driver (Module) | N/A | isci |
| Controller Driver Version | N/A | 1.1.0-rh |
| Controller firmware | N/A | SCU 3.8.0.1029 |
| RAID configuration | N/A | None |
| USB ports | | |
| Number | N/A | 4 |
| Type | N/A | 2.0 |

Figure 2: System configuration information for the test systems.

APPENDIX C – DETAILED TEST METHODOLOGY

Setting up and configuring the HP ProLiant m700 servers in the Moonshot system

We set up two auxiliary servers to support PXE booting the AMD-based HP ProLiant m700 servers: the first ran CentOS 6.5 and provided NFS storage for the nodes' root directories, and the second provided NTP, DNS, DHCP, and TFTP services to supply each node with an IP address, boot image, and path to its root directory.

Configuring the Moonshot Chassis Management (CM) and 180G Switch modules

1. Log onto the Moonshot CM via its serial interface as administrator.
2. Set its networking settings, IP address, mask, gateway, and DNS and NTP servers, as in the following commands:

```
set network ip 10.10.10.4
set network mask 255.255.255.0
set network gateway none
set network dns 1 10.10.10.10
set ntp primary 10.10.10.10
disable winsreg
disable ddnsreg
```

3. Reset the CM to effect these changes:

```
reset cm
```

4. Connect to the CM via ssh and log on as administrator.
5. Print the MAC addresses of the node's Ethernet interfaces:

```
show node macaddr all
```

6. Capture these from the console screen (e.g., by selecting with the mouse and copying), and save them to a file on the PXE server for use in the next section. The output will resemble the following:

```
Slot ID      NIC 1 (Switch A)  NIC 2 (Switch B)  NIC 3 (Switch A)  NIC 4
(Switch B)
-----
1  c1n1  2c:59:e5:3d:3e:a8 2c:59:e5:3d:3e:a9 N/A              N/A
1  c1n2  2c:59:e5:3d:3e:aa 2c:59:e5:3d:3e:ab N/A              N/A
1  c1n3  2c:59:e5:3d:3e:ac 2c:59:e5:3d:3e:ad N/A              N/A
1  c1n4  2c:59:e5:3d:3e:ae 2c:59:e5:3d:3e:af N/A              N/A
```

7. Connect to the Moonshot 180G Switch module:

```
connect switch vsp all
```

8. Log onto the switch as admin.

9. Enter privilege mode:

```
enable
```

10. Set the switch's IP address:

```
serviceport protocol none
```

```
serviceport ip 10.10.10.3 255.255.255.0
```

11. Enter global configuration mode:

```
configure
```

12. Set the second 40Gbps QSFP+ port to run in 4x10G mode:

```
interface 1/1/6
hardware profile portmode 4x10g
ctrl-z
write memory
reload
```

13. Activate all ports:

```
shutdown all
no shutdown all
```

14. Exit the privileges modes by press Ctrl-Z twice.

15. Log off the switch by typing `quit`

16. When prompted, type `y` to save the configuration.

17. Exit the switch module console and return to the CM console by pressing ESC.

Configuring the auxiliary PXE and NFS servers for diskless ProLiant m700 servers

We configured the auxiliary NFS server (DNS name `NFS_SERVER`) to export directory `NFS_PATH` to the nodes' subnet (10.10.10.0/24) and created root directories for each node using the naming convention: `c01n1`, `c01n2`, `c01n3`, `c0n4`, `c02n1`, ..., `c45n4`. The second server provided the following services to the ProLiant m700 nodes:

1. DNS resolution of the nodes' hostnames. The following excerpt is from the file `/etc/hosts`.

```
10.10.10.51 c01n1
10.10.10.52 c01n2
10.10.10.53 c01n3
10.10.10.54 c01n4
```

2. DHCP service provides each node with an IP address, netmask, DNS server, NTP server, name (common) boot image, and the address of the TFTP server to obtain this image. The following excerpt is from the file `/etc/dhcp/dhcpd.conf` and shows the global DHCP configuration.

```
allow booting;
allow bootp;
log-facility local7;

option subnet-mask          255.255.255.0;
option broadcast-address    10.10.10.255;
option domain-name-servers 10.10.10.10;
```

```
option ntp-servers          10.10.10.10;
option time-offset         -5;
```

3. We used a simple awk script to parse the contents of the file of MAC address from step 6 in the previous section and to create node-specific DHCP entries in `/etc/dhcp/dhcpd.conf`. We used the following template for the DHCP entry for each node (replacing `FIX_HOSTNAME`, `FIX_HOST_MAC`, and `FIX_HOST_IP` in the template with the correct values for the node):

```
group {
    filename      "/pxelinux.0";
    next-server   10.10.10.10;

    host FIX_HOSTNAME {
        hardware ethernet FIX_HOST_MAC;
        fixed-address FIX_HOST_IP;
    }
    ...
}
```

4. TFTP service provides boot images and root-directory location to each node. Create the directories `/var/lib/tftp/centos6` and `/var/lib/tftpboot/pxelinux.cfg`:

```
mkdir /var/lib/tftp/centos6 /var/lib/tftpboot/pxelinux.cfg
```

5. Copy the PXE file to `/var/lib/tftp`, and the OS images to `/var/lib/tftp/centos6`:

```
cp /usr/share/syslinux/pxelinux.0 /var/lib/tftpboot
cp /boot/initramfs-2.6.32-431.11.2.el6.x86_64.img /var/lib/tftp/centos6
cp vmlinuz-2.6.32-431.11.2.el6.x86_64 /var/lib/tftp/centos6
```

6. We used a simple awk script to parse the contents of the file of MAC address from step 6 in the previous section and to create node-specific PXE files in directory `/var/lib/tftpboot/pxelinux.cfg/`. The name of a node's PXE file is "01-" followed by the node's MAC address in hexadecimal with hyphens between pairs of characters; for example, `01-2c-59-e5-3d-3e-a8`. We used the following template to create the contents of each file (replacing `FIX_HOSTNAME` in the template with the correct values for the node). Again, `NFS_SERVER:/NFS_PATH` is to be replaced with the NFS handle for the share containing the nodes' root directories. The template contains the following:

```
default linux
prompt 0
serial 0 9800n8
label linux
    kernel centos6/vmlinuz-2.6.32-431.11.2.el6.x86_64
```

```

    append initrd=centos6/initramfs-2.6.32-431.11.2.el6.x86_64.img
    console=tty0 \
    console=ttyS0,9600n8 root=nfs:NFS_SERVER:/NFS_PATH/FIX_HOSTNAME rw
    ip=dhcp

```

7. Change the permissions of the tftp directory and files beneath it so that all users can read them.

```
chmod -R a+rX /var/lib/tftp
```

Installing and configuring the operating system *en masse*

1. Log onto the CentOS auxiliary server (PXE server) as root.
2. Mount the NFS directory for nodes' root directories with the rootsquash option at mountpoint /opt/diskless.
3. Create a list of node names:

```

echo
c{0{1,2,3,4,5,6,7,8,9},{1,2,3}{0,1,2,3,4,5,6,7,8,9},4{0,1,2,3,4,5}}n{1,2,3,4
}\
    > /opt/nodes.txt

```

4. Create the root directory for each node:

```

for node in $(cat /opt/nodes.txt); do
    mkdir /opt/diskless/${node}
done
chmod -R a+rx /opt/diskless

```

5. Install the CentOS base package group and the following miscellaneous package group on all the nodes:

```

for node in $(cat /opt/nodes.txt); do
    yum --installroot=/opt/diskless/${node} install -y @base @compat-libraries
\
    @console-internet @fonts @hardware-monitoring @large-systems @legacy-unix
\
    @legacy-x @network-tools @performance @perl-runtime @system-admin-tools
done

```

6. Set the hostname of each node and disable SELinux:

```

for node in $(cat /opt/nodes.txt); do
    echo "HOSTNAME=${node}" > /etc/sysconfig/network
    echo NETWORKING=yes >> /etc/sysconfig/network
    sed -i 's/^SELINUX=enabled/SELINUX=disabled/' /etc/selinux/config
done

```

Booting the HP ProLiant m700 servers

1. Power on the PXE and NFS auxiliary servers.
2. Log onto the Moonshot CM as administrator.

3. Power on every node:

```
set node power on all
```

Installing the AMD OpenCL libraries

1. Download the AMD Catalyst 14.10.1006-1 drivers for 64-bit Linux, and copy the installer to the PXE server.
2. Log onto the PXE server as root.
3. Uncompress and change the execution permissions of the AMD Catalysts installer:

```
unzip amd-catalyst-14.1-betav1.3-linux-x86.x86_64.zip
chmod a+rx amd-driver-installer-13.35.1005-x86.x86_64.run
```
4. Build an RPM package for the Catalyst software:

```
./amd-driver-installer-13.35.1005-x86.x86_64.run --buildpkg RedHat/RHEL6_64a
```
5. Install the Catalyst software on the live nodes:

```
for node in $(cat /opt/nodes.txt) ; do
  scp fglrx64_p_i_c-14.10.1006-1.x86_64.rpm ${node}:/tmp/
  ssh ${node} yum localinstall -y /tmp/fglrx64_p_i_c-14.10.1006-1.x86_64.rpm
done
```

Setting up and configuring the Intel server

Configuring disk volumes and BIOS

1. From the RAID-controller configuration page, connect two disks as a RAID 1 volume.
2. From the BIOS configuration screen, reset all settings to their default values.
3. Set the server power configuration to maximum performance.

Installing the CentOS 6.5 64-bit operating system

1. Insert the CentOS 6.5 installation DVD and boot from it.
2. On the Welcome to CentOS 6.5! screen, select Install or upgrade an existing system, and press Enter.
3. On the Disc Found screen, select Skip, and press Enter.
4. On the CentOS 6 screen, click Next.
5. On the installation-selection screen, keep the default, and click Next.
6. On the keyboard-selection screen, keep the default, and click Next.
7. On the storage-selection screen, click Basic Storage Devices, and click Next.
8. On the Storage Device Warning pop-up screen, click Yes, discard any data.
9. On the Hostname screen, enter the servers name and click Configure Network.
10. On the Network Connections pop-up screen, click Add.
11. On the Choose a Connection Type selected Wired, and click Create.
12. On the Editing Wired Connection pop-up, select the IPv4 Settings tab, change Method to Manual, click Add, enter the interface's IP address, netmask and gateway, and click Apply.
13. Close the Network Connections pop-up screen.

14. Click next on the Hostname screen.
15. On the time-zone screen, click Next.
16. On the administrator-password screen, enter the Root Password (twice), and click Next.
17. On the Which type of installation would you like screen, click both Replace Existing Linux Systems(s), and click Next.
18. On the Format Warnings pop-up screen, click Format.
19. On the Writing storage configuration to disk pop-up screen, click Write changes to disk.
20. On the boot-loader selection screen, click Next.
21. On the software-selection screen, click Basic Server, and click Next.
22. On the Congratulations screen, click Reboot.

Configuring the operating system

1. Log onto the server as root.
2. Set the hostname.
3. Install additional system software:


```
for node in $(cat /opt/nodes.txt); do
    yum --installroot=/opt/diskless/${node} install -y @base @compat-libraries
    \
    @console-internet @fonts @hardware-monitoring @large-systems @legacy-unix
    \
    @legacy-x @network-tools @performance @perl-runtime @system-admin-tools
done
```
4. Disable SELinux:


```
sed -i 's/^SELINUX=enabled/SELINUX=disabled/' /etc/selinux/config
```
5. Reboot the server:


```
shutdown -r now
```

Installing the Intel OpenCL libraries

1. Download the Intel OpenCL SDK for 64-bit Linux, version XE 2013 R3, and copy it onto the Intel white-box server.
2. Log onto the Intel white-box server as root.
3. Extract the software:


```
tar zxf
intel_sdk_for_ocl_applications_xe_2013_r3_runtime_3.2.1.16712_x64.tgz
```
4. Import the Intel signing:


```
rpm --import Intel-E901-172E-EF96-900F-B8E1-4184-D7BE-0E73-F789-186F.pub
```
5. Install the RPM:


```
cd intel_sdk_for_ocl_applications_xe_2013_r3_runtime_3.2.1.16712_x64
```

```
yum localinstall opencl-1.2-base-3.2.1.16712-1.x86_64.rpm \  
opencl-1.2-intel-cpu-3.2.1.16712-1.x86_64.rpm
```

Installing the rendering software on the ProLiant m700 and white-box servers

1. Download the LuxRender software, including the Blender plugin, from www.luxrender.net as lux-v1.3.1-x86_64-sse2-OpenCL.tar.bz2.

2. Download the workload, a sample scene, from 3developer.com/sala/sala-lux.zip.

3. Copy the software and workload to each server or node and extract the files. For example,

```
scp lux-v1.3.1-x86_64-sse2-OpenCL.tar.bz2 sala-lux.zip c01n1:  
ssh c01n1 tar jxf lux-v1.3.1-x86_64-sse2-OpenCL.tar.bz2  
ssh c01n1 unzip sala-lux.zip
```

4. Copy the following test harness, fg.sh, to each node and the white-box server:

```
#!/bin/bash  
# test harnerss  
pkill luxconsole  
R=0203  
NUM=2  
S=/root/sala/Sala.blend.lxs  
F=RUN_$(hostname)  
sync  
echo 3 > /proc/sys/vm/drop_caches  
rm -rf /root/sala-2/Sala.Scene.*.flm 2>&1 > /dev/null  
rm -rf /tmp/cache-11_*/*.flm 2>&1 > /dev/null  
for i in $(seq $NUM) ; do  
    echo $i  
    tag=/tmp/cache-11_$i  
    mkdir $tag 2>&1 > /dev/null  
    /root/lux-v1.3.1-x86_64-sse2-OpenCL/luxconsole -o $tag/out-11-$i $$ \  
        &> $tag/${F}_$i-$R.txt &  
done#
```

5. To start the workload on the nodes, run the following commands from the PXE server:

```
for node in $(cat /opt/nodes.txt) ; do  
    echo $node  
    ssh $node sh fg.sh  
done
```

6. To run the workload on the white-box server, run the following command from the PXE server, where IP_WHITEBOX is the IP address for the white-box server. When the operation is finished, the computation rate is stored in the directories /tmp/cache*.

```
ssh IP_WHITEBOX sh fg.sh
```

Measuring power

To record each server's power consumption during each test, we used five Extech Instruments (www.extech.com) 380803 Power Analyzer/Dataloggers. We connected each power cord from the servers under test to its own Power Analyzer output-load power-outlet. We then plugged the power cord from each Power Analyzer's input voltage connection into a power outlet.

We used the Power Analyzer's Data Acquisition Software (version 3.0) to capture all recordings. We installed the software on a separate PC, which we connected to the Power Analyzer via an RS-232 cable. We captured power consumption at one-second intervals.

We then recorded the power usage (in watts) for each system during the testing at one-second intervals. To compute the net usage, we averaged the power usage during the time the system was producing its peak performance results. We call this time the power measurement interval.

ABOUT PRINCIPLED TECHNOLOGIES



Principled Technologies, Inc.
1007 Slater Road, Suite 300
Durham, NC, 27703
www.principledtechnologies.com

We provide industry-leading technology assessment and fact-based marketing services. We bring to every assignment extensive experience with and expertise in all aspects of technology testing and analysis, from researching new technologies, to developing new methodologies, to testing with existing and new tools.

When the assessment is complete, we know how to present the results to a broad range of target audiences. We provide our clients with the materials they need, from market-focused data to use in their own collateral to custom sales aids, such as test reports, performance assessments, and white papers. Every document reflects the results of our trusted independent analysis.

We provide customized services that focus on our clients' individual requirements. Whether the technology involves hardware, software, Web sites, or services, we offer the experience, expertise, and tools to help our clients assess how it will fare against its competition, its performance, its market readiness, and its quality and reliability.

Our founders, Mark L. Van Name and Bill Catchings, have worked together in technology assessment for over 20 years. As journalists, they published over a thousand articles on a wide array of technology subjects. They created and led the Ziff-Davis Benchmark Operation, which developed such industry-standard benchmarks as Ziff Davis Media's Winstone and WebBench. They founded and led eTesting Labs, and after the acquisition of that company by Lionbridge Technologies were the head and CTO of VeriTest.

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.
