



EQUALLOGIC

TECHNICAL SUPPORT INSTRUCTIONAL SERIES

Technical Instruction	Document ID: N/A
Installing and Using IOMeter (Dynamo) Under Linux	Revision: A02
	Revision Date: September 21, 2009

Installing and Using IOMeter Under Linux

Introduction

IOMeter allows you to test the performance of the network and your PS array for iSCSI traffic. This document assumes you are familiar with Linux, have properly installed the iSCSI initiator on your Linux host, and can connect to iSCSI targets on your PS array. IOMeter performs a configurable suite of I/O operations on a volume on your array and measures the performance. This test verifies your network configuration for optimal connectivity with your PS Array.

On Linux, IOMeter is composed to two separate pieces: the host process called *dynamo*, which is responsible for performing the actual I/O operations; and the graphical IOMeter client which communicates with *dynamo* and presents the results. *dynamo* must be run on the host which is connected to the iSCSI volume, but the graphical IOMeter client only exists for Windows. So you must have a Windows machine available with network access to the Linux host running *dynamo*. You must have root access on your Linux host to perform these actions.

Finally, under Linux, the kernel will normally cache storage accesses, so to test real I/O performance with a PS array, you have to use an uncached, raw device to access the system. One way to achieve this is through the "raw" virtual device program, which is what we will describe below.

Basic Steps:

1. Download and/or compile *dynamo*
2. Download and install IOMeter for Windows
3. Create a volume on your array and connect it to your host
4. Configure raw device
5. Start *dynamo*
6. Start IOMeter client, configure Disk Target
7. Configure Access Specifications
8. Run Tests

Download and Compile *dynamo*

The IOMeter project lives at SourceForge, and can be found here:

<http://sourceforge.net/projects/iometer/>

Click on the Download IOMeter link, and select the "iometer-stable" package from 2006-07-27. Download the most recent source package. There are five files associated with the IOMeter project:

```
iometer-2006_07_27.common-docu.zip
iometer-2006_07_27.common-src.zip
iometer-2006_07_27.linux.i386-bin.tgz
iometer-2006_07_27.netware.i386-bin.zip
iometer-2006.07.27.win32.i386-setup.exe
```

For Linux, you may be able to work with the pre-compiled version in the `linux.i386-bin.tgz` file, but otherwise, you can use the `common-src.zip` file to build it yourself. The `common-docu.zip` file may also be useful for you. For now, we will assume that you have to build your own copy. Download the `iometer-2006_07_27.common-src.zip` file.

Untar the package, and `cd` into the “src” directory.

Then build the package. Run the following commands:

```
# make -f Makefile-Linux.i386 clean
# make -f Makefile-Linux.i386 all
/usr/bin/g++ -O -Wall -D_GNU_SOURCE -DIOMTR_OSFAMILY_UNIX -DIOMTR_OS_LINUX
-DIOMTR_CPU_X86_64 -DIOMTR_SETTING_GCC_M64 -D_LARGEFILE_SOURCE -
D_LARGEFILE64_SOURCE -D_FILE_OFFSET_BITS=64 -D_REENTRANT -DNOMINMAX -c
IOGlobals.cpp
...
```

You may also choose to use `Makefile-Linux.i386`. This will give you the *dynamo* binary in the current directory. You can simply run the program from there without having to install it anywhere else.

Download and Install IOMeter for Windows

The IOMeter project lives at SourceForge, and can be found here:

<http://sourceforge.net/projects/iometer/>

Click on the Download IOMeter link, and select the "iometer-stable" package from 2006-07-27. Download and run the self-installing Windows host and client package on your chosen Windows machine. This will install IOMeter in your Start Menu for easy access later.

```
iometer-2006.07.27.win32.i386-setup.exe
```

Create Volume

Create a volume on your PS array, then connect to it via iSCSI. The volume does not need to be large: 5GB is reasonable. There is no need to initialize or otherwise manipulate the volume; IOMeter uses it as a raw device. Do not put a filesystem on it.

You must create and attach to your volume before *dynamo* starts. *dynamo* doesn't refresh its view of the local drive table without being restarted. Note that you may need to reload your iSCSI daemon after creating the new volume, to get it attached to your Linux host. Also, make sure you know which physical device your new volume corresponds to.

Configure Raw Device

In order to provide a real network performance metric, we need to disable local I/O caching for the iSCSI device. We do that by mapping a “raw” device to the SCSI disk associated with the volume. “`man raw`” will give you more information on the process.

We must identify the SCSI drive we wish to use, and then we have to check the existing raw devices to determine whether this drive has already been mapped, and if not, what the first available raw device is. Use “iscsi-ls” and “iscsi-device” to positively identify which SCSI device corresponds to your selected volume. Then run “ls -l <SCSI-device>” to see the major and minor device numbers. Finally, run “raw -qa” to show the complete list of configured raw devices.

If any configured raw device is mapped to the major and minor numbers of your volume, use that raw device. Otherwise, choose a new device other than the configured ones, and assign your SCSI device to it using this command (assuming raw3 is the next available device and /dev/sdc is your iSCSI volume): “raw /dev/raw/raw3 /dev/sdc”.

Start *dynamo*

You can start *dynamo* and the IOMeter graphical client in any order, but in this example we will first start *dynamo*. You don’t necessarily have to be root to run *dynamo*, but you need to have been root to configure the raw devices, and there may be some statistical data you cannot access as a non-privileged user. And your list of devices might not be complete. So we recommend staying root.

To run *dynamo*:

```
# ./dynamo -i <Windows IP> -m <Linux IP> -n <display name> -d <raw device>
```

You can also specify host names instead of IP addresses, but if you’re at all unsure of your DNS coverage, you might as well use the IP address. There is generally little risk of harm in using incorrect values, but you will not get useful results. <Linux IP> is the address of the host running *dynamo*; <Windows IP> is the address of the machine running the IOMeter graphical client; and <display name> is simply the machine name in the client that the Linux host will appear as. If you omit the display name, the *dynamo* defaults to the hostname. <raw device> is the raw device you created above, usually “/dev/raw/raw1”.

Start IOMeter

On your Windows machine, IOMeter will have created a program group for itself. Start it from that.

Under “Topology”, open the system tree and select the heading that corresponds to the display name you set. Make sure only one worker (“Worker 1”) is configured for that host. (If there is more than one worker, remove all but Worker 1 by selecting the others and clicking on the “Disconnect Selected Worker or Manager” icon; it’s supposed to look like data cables being pulled apart, with arrows pointing outward.)

Click on “Worker 1”, and note that the Disk Targets tab populates with the available devices from your Linux host. Note that disks with filesystems have red slashes through them, which is IOMeter’s way of discouraging you from using them. Select the raw device you created. Then set the “Maximum Disk Size” field to 8192 and the “# of Outstanding I/Os” to 64.

Note - “Maximum Disk Size” of 8192 sectors is to arbitrarily limit it, so nearly all I/Os will be “out of the cache” after the first set. Therefore, the array will turn I/Os around as quickly as they come in. This test is mainly to focus on verifying the switch/NIC/HBA configuration is correct.

Configure Access Specifications

Access Specifications define what kind of I/O is going to be performed. You can define the data access profile with mixtures of reads and writes, random and sequential accesses, burst characteristics, and other options. We recommend using the predefined specifications until you are familiar enough to want to make more specific tests.

Select the “Access Specifications” tab. There are two columns: Global Access Specifications and Assigned Access Specifications. Choose one from the Global column, and hit the “<< Add” button to place it in the Assigned column. If you wish to change your access specifications, you must first select the current Assigned Access Specification, and “Remove >>” it. Then select another Access Specification and “<< Add” it. You should only have one Assigned Access Specification at a time.

We recommend testing initially with two of the Global Specifications: “32K; 0% Read; 0% Random” will test large sequential writes; and “32K; 100% Read; 0% Random” will test large sequential reads.

Run Tests

Finally, select the “Results Display” tab. Slide the “Update Frequency” slider to 2 seconds. In the “Results Since” groupbox, select the radio button for “Last Update”. Go to the last tab, Test Setup, and change the Record Results pulldown in the middle of the tab to None.

Then, go back to the “Results Display” tab, and click on the green flag button at the top of the Iometer window to start the test. The test will now start running. Wait for a few updates, and you should see >60 MB/sec throughput in the 0% Read mode.

After about 20 seconds of fairly stable results, you may stop the test and select another Access Specification, and then start the tests again. Wait for a few updates, and you should again see >60 MB/sec throughput in the 100% Read mode as well.

Document Control and History Information

Revision	Change Author/Organization	Reason	Approved By/Organization	Date
A02	M. Cynewski / GCSS	ISO Compliance	M. Gauthier / GCSS	9/21/2009
A01	Unknown	Initial Release	Unknown	8/11/2007