

Online Lab Guide

Sun Storage 7000 Unified Storage Systems

version 7, 3 November 2009



Part III – Analytics

3. Analytics

In this section of the lab, you'll work with the Sun Storage Unified Storage System's unique analytical features. This storage system is equipped with an advanced DTrace-based facility for server analytics. Analytics provides real time graphs of various statistics, which can be saved for later viewing. About a dozen high level statistics are provided which can then be customized to provide lower level details. Groups of viewed statistics can be saved as worksheets for future reference.

Specifically in this part of the lab, you will ...

- initiate some scripts which will generate artificial load on the storage system
- set up some monitoring worksheets
- use analytics to drill down on the worksheet diagnostics

You'll do this for both the iSCSI attached storage and the CIFS attached storage.

The setup and configuration sections of the lab that you have completed initialized the storage system, and then created and shared some slices of the storage using 3 different protocols, and these slices were tied into the Windows client in your lab workspace.

A. Overview of Analytics

Analytics is an advanced facility to graph a variety of statistics in real-time and record this data for later viewing. It has been designed for both long term monitoring and short term analysis. When needed, it makes use of DTrace to dynamically create custom statistics, which allows different layers of the operating system stack to be analyzed in detail.

Analytics has four main features

1. Drilldown Analysis

Analytics has been designed around an effective performance analysis technique called drill-down analysis. This involves checking high level statistics first, and to focus on finer details based on findings so far. This quickly narrows the focus to the most likely areas.

2. Statistics

In Analytics, the user picks statistics of interest to display on custom worksheets. Statistics available from Analytics include:

- Network device bytes by device and direction
- NFS operations by filename, client, share, type, offset, size and latency
- CIFS operations by filename, client, share, type, offset, size and latency
- Disk operations by type, disk, offset, size and latency
- CPU utilization by CPU-id, mode and application

The Open Worksheets view shows a listing statistics, and the Preferences view allows you to enable advanced Analytics which will make

many more statistics available. The Statistics view discusses available statistics in more detail.

3. Datasets

A dataset refers to all existing data for a particular statistic. Datasets contain:

- statistic data cached in memory due to the statistic being opened or archived
- archived statistic data on disk

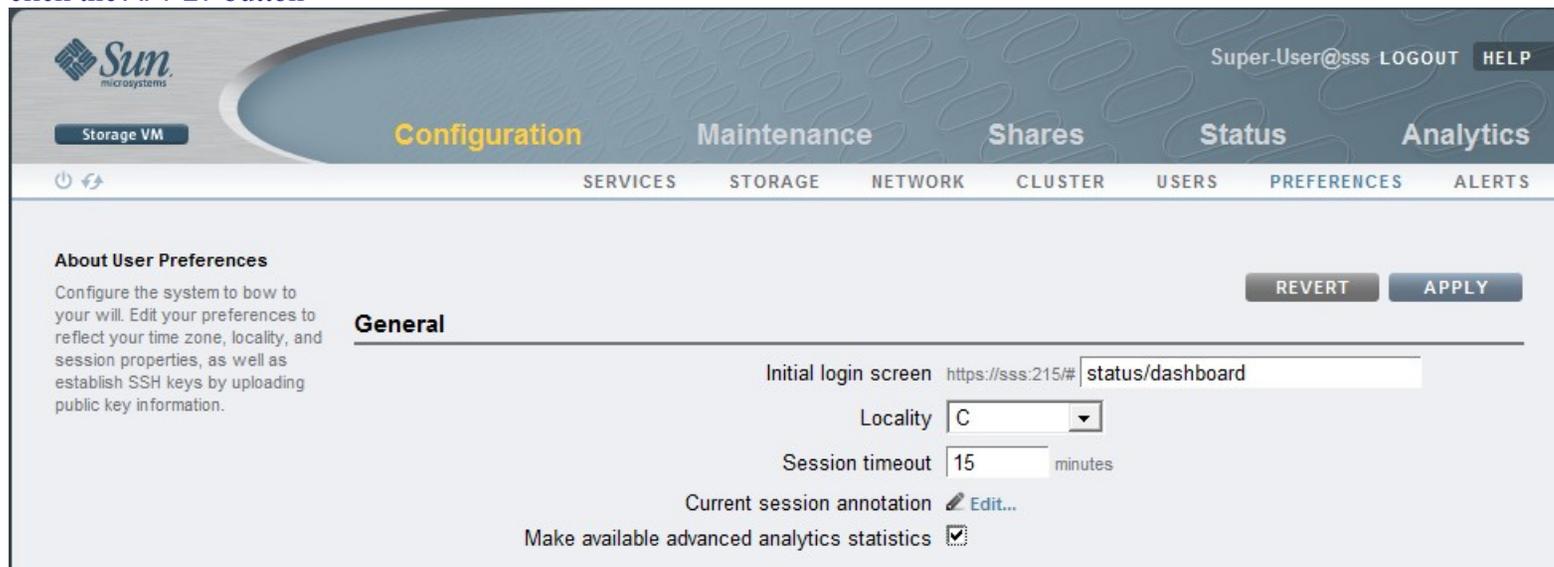
Datasets can be managed in the Datasets view.

4. Worksheets

A worksheet is the BUI screen on which statistics are graphed. Multiple statistics can be plotted at the same time, and worksheets may be assigned a title and saved for future viewing. The act of saving a worksheet will automatically execute the archive action on all open statistics. This means that whatever statistics were open, will continue to be read and archived forever until you make a change.

By default, the Statistics available under Analytics are basic and useful to most storage administrators. But to demonstrate the power of Analytics, you need to turn on all of the available features.

1. on the menu, select Configuration → Preferences
2. select the checkbox labeled Make available advanced analytics statistics
3. click the APPLY button



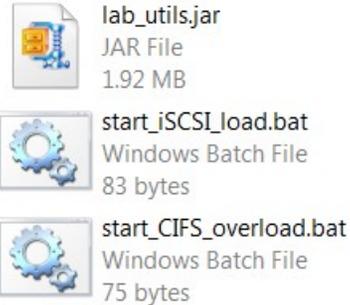
The screenshot shows the Sun Storage VM web interface. At the top, there is a navigation bar with the Sun logo, a 'Storage VM' button, and tabs for 'Configuration', 'Maintenance', 'Shares', 'Status', and 'Analytics'. Below this is a secondary navigation bar with 'SERVICES', 'STORAGE', 'NETWORK', 'CLUSTER', 'USERS', 'PREFERENCES', and 'ALERTS'. The main content area is titled 'About User Preferences' and contains a 'General' tab. Under the 'General' tab, there are several settings: 'Initial login screen' with a text input field containing 'https://sss:215/#status/dashboard', 'Locality' with a dropdown menu set to 'C', 'Session timeout' with a text input field set to '15' and the unit 'minutes', 'Current session annotation' with an 'Edit...' link, and 'Make available advanced analytics statistics' with a checked checkbox. There are 'REVERT' and 'APPLY' buttons in the top right corner of the settings area.

B. iSCSI Attached Storage

1. Generate Load

First, we need to create some load on the storage system (files being written back and forth to the allocated storage).

1. go to the Windows desktop in your lab environment and open (double-click on) a folder named lab_utils
2. inside this folder should be 3 files



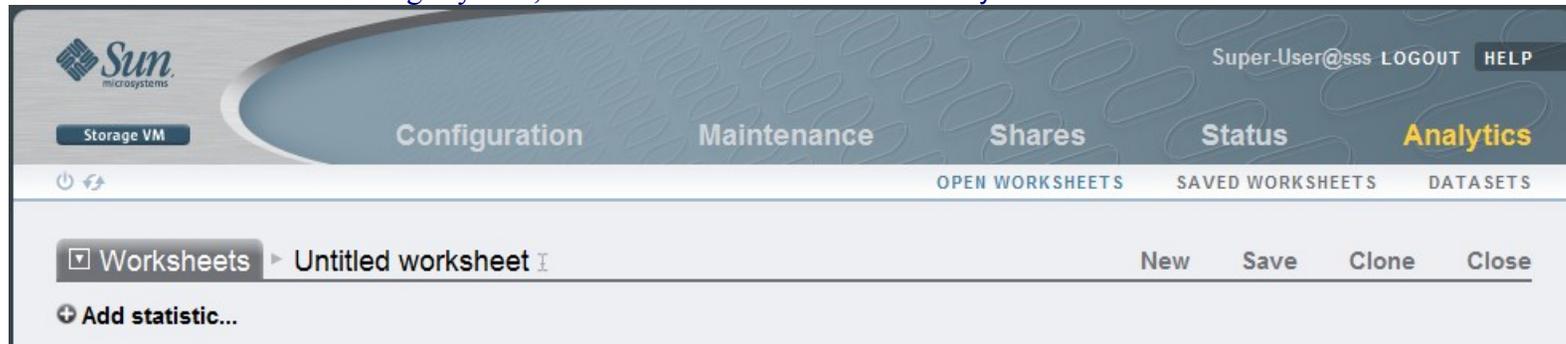
3. open or start (by double-clicking on it) the file named start_iSCSI_load.bat

Your Windows client is now copying a large image back and forth to the the storage system.

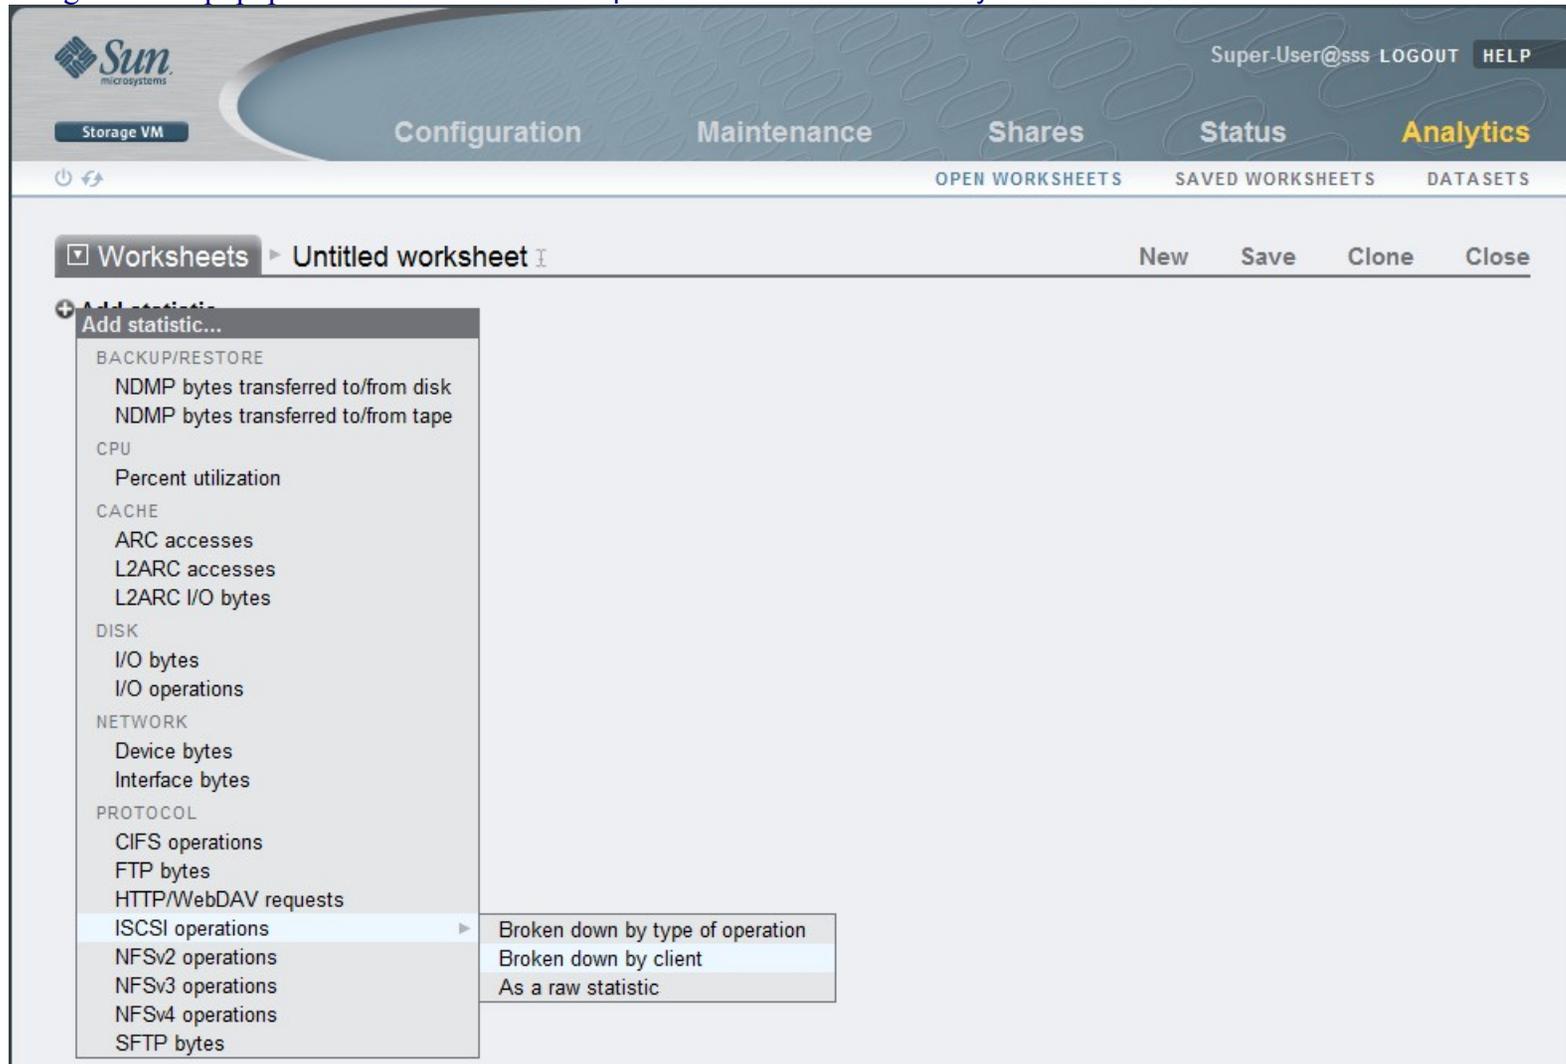
2. Setup Worksheets

Second, you will want to setup an Analytics worksheet to observe the storage system's behavior under the load being generated on it.

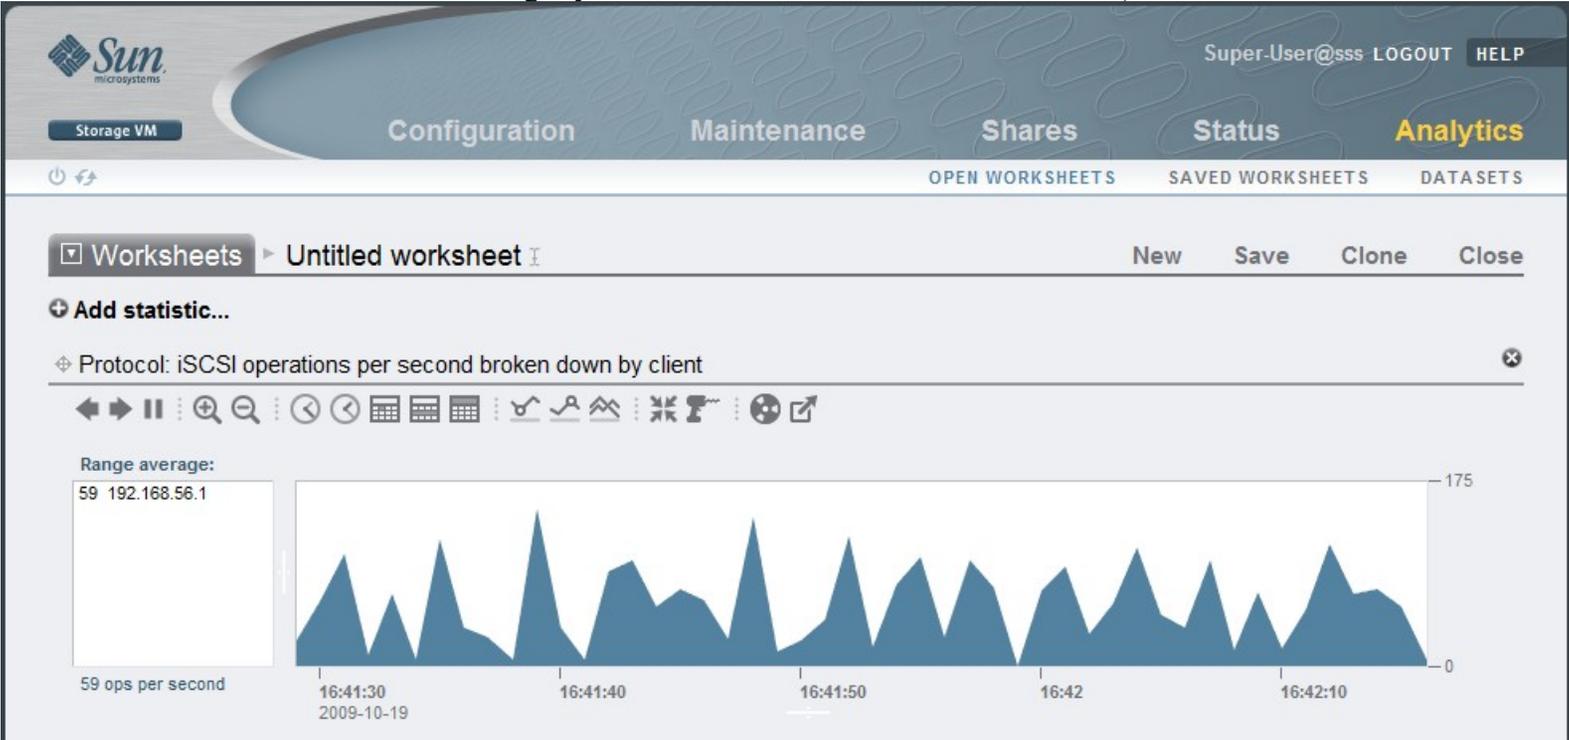
1. in the Firefox BUI for the storage system, click on the menu item titled Analytics



2. you get an empty untitled worksheet
3. start to add a Statistic by clicking on the Add statistic ... label
4. the popup menu shows you the high-level categories for available statistics
5. navigate in the popup menu to the item iSCSI operations → Broken down by client and select it



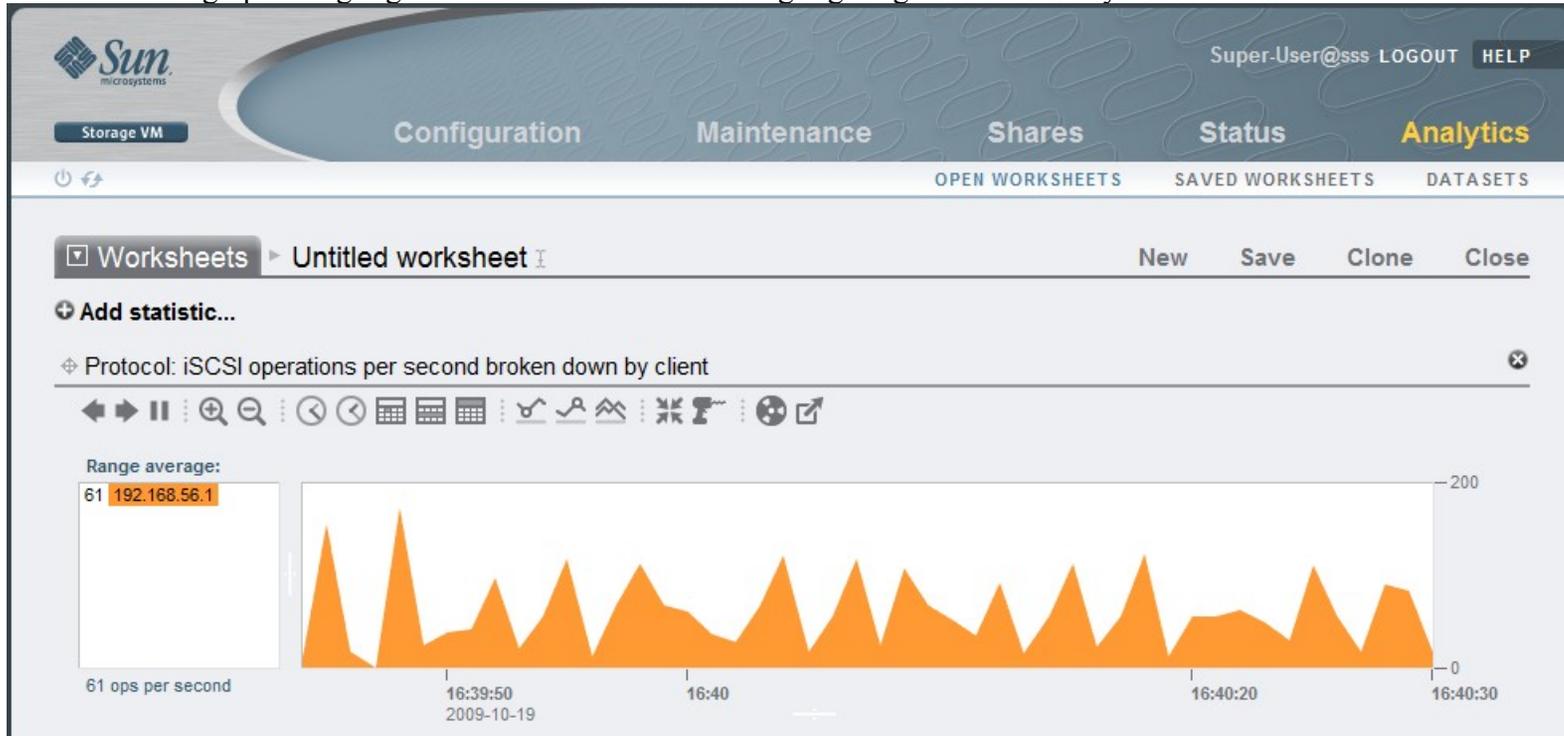
6. a worksheet will appear showing you the traffic being generated by client IP address (note that you're seeing an IP address instead of a DNS name for the client since the storage system has not been tied into a DNS service)



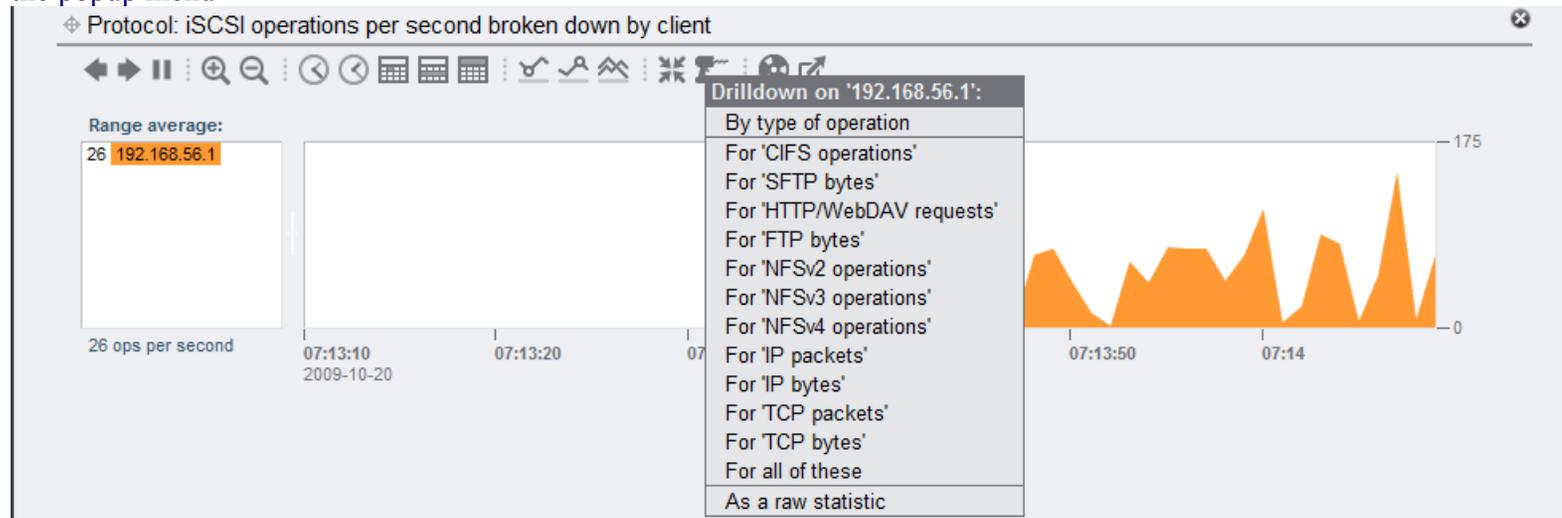
3. Drill Down

From this initial statistic, we can drill down further to examine the behavior.

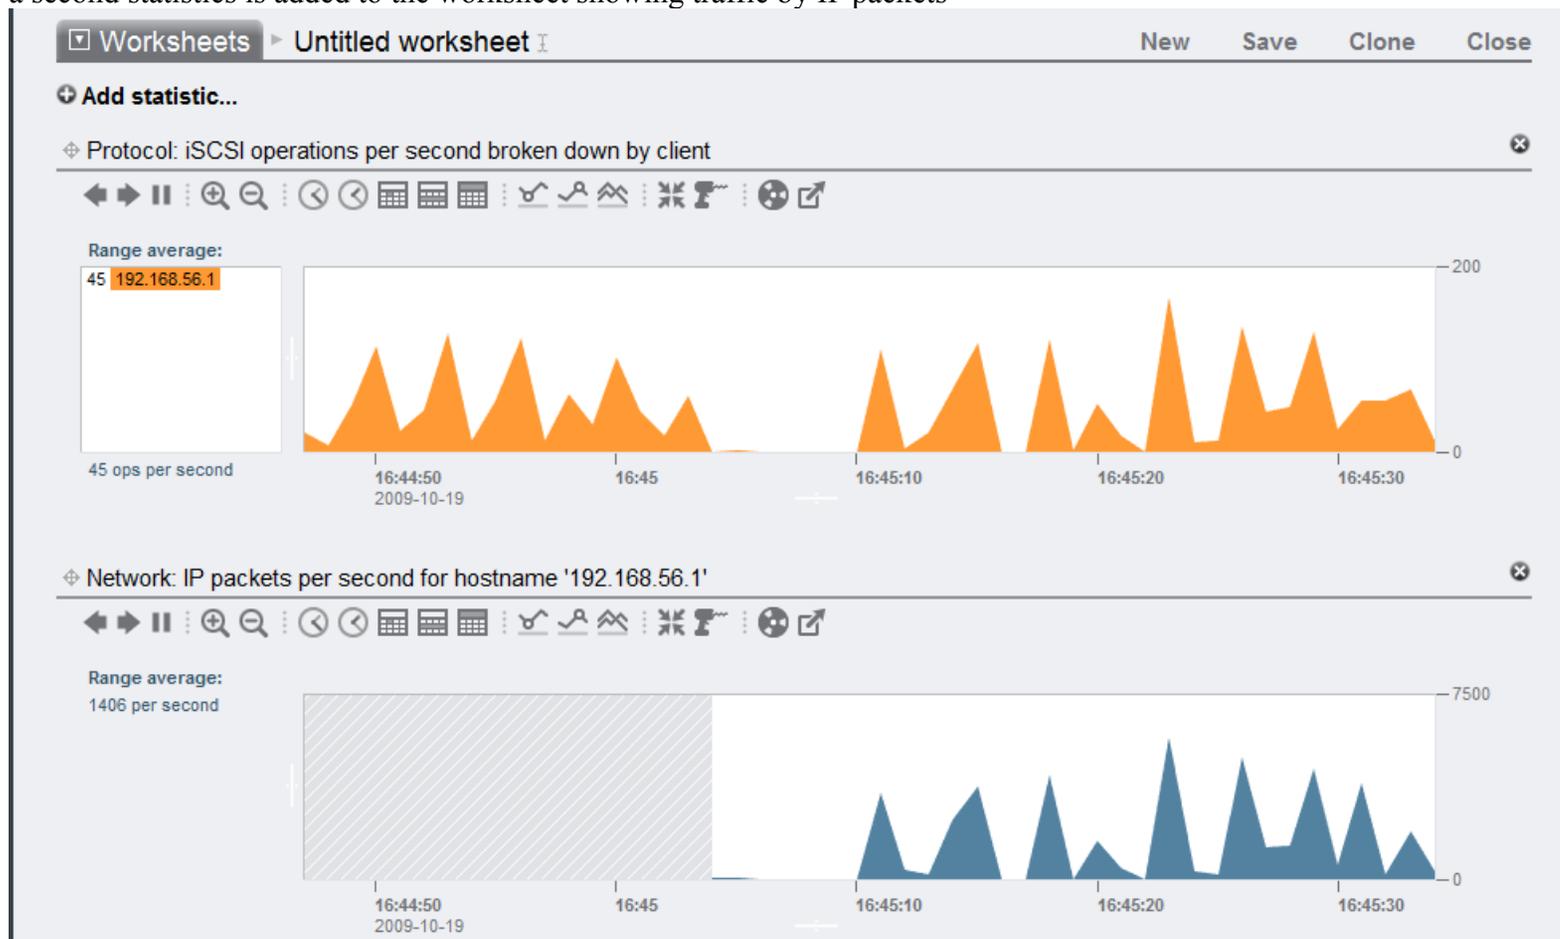
1. in order to drill down and see more information, you need to select a particular client, so [click on the IP address in the Range average list](#)
2. notice how the graph is highlighted in the same color now highlighting the IP address you selected



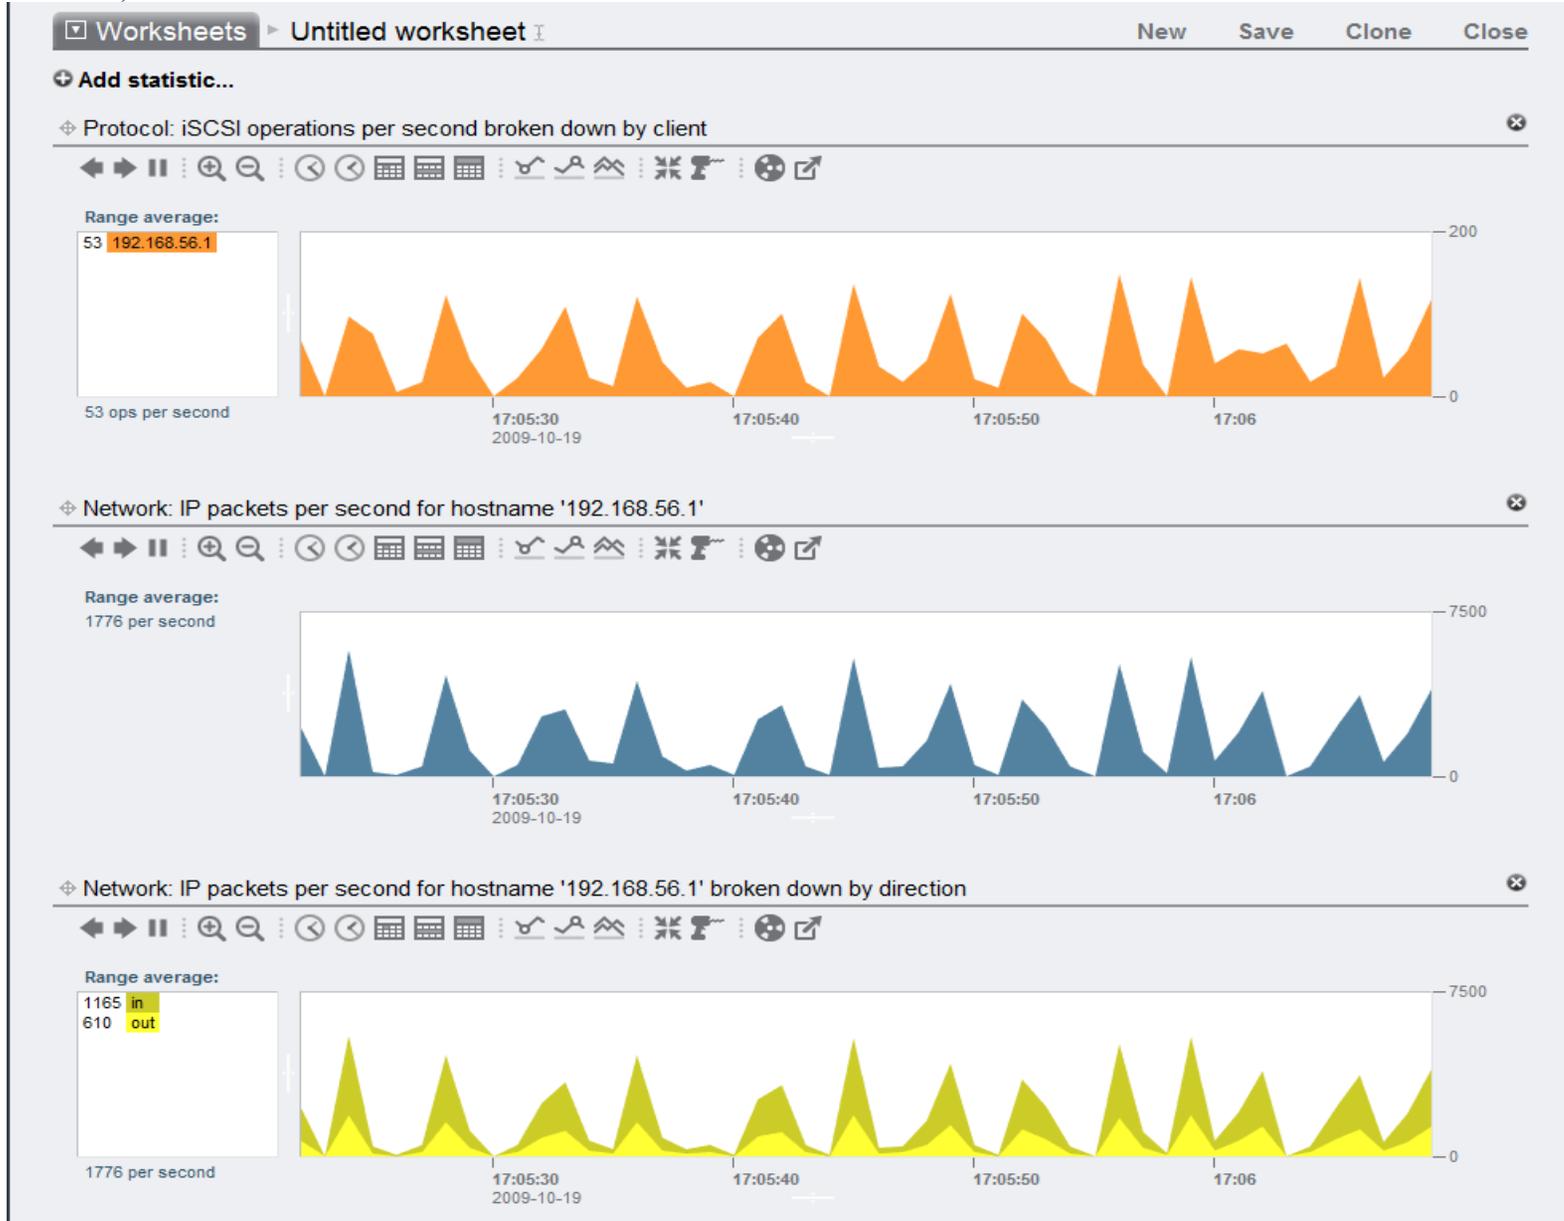
3. now it gets interesting – click on the power drill icon on the toolbar for the currently shown statistic and select For 'IP packets' from the popup menu



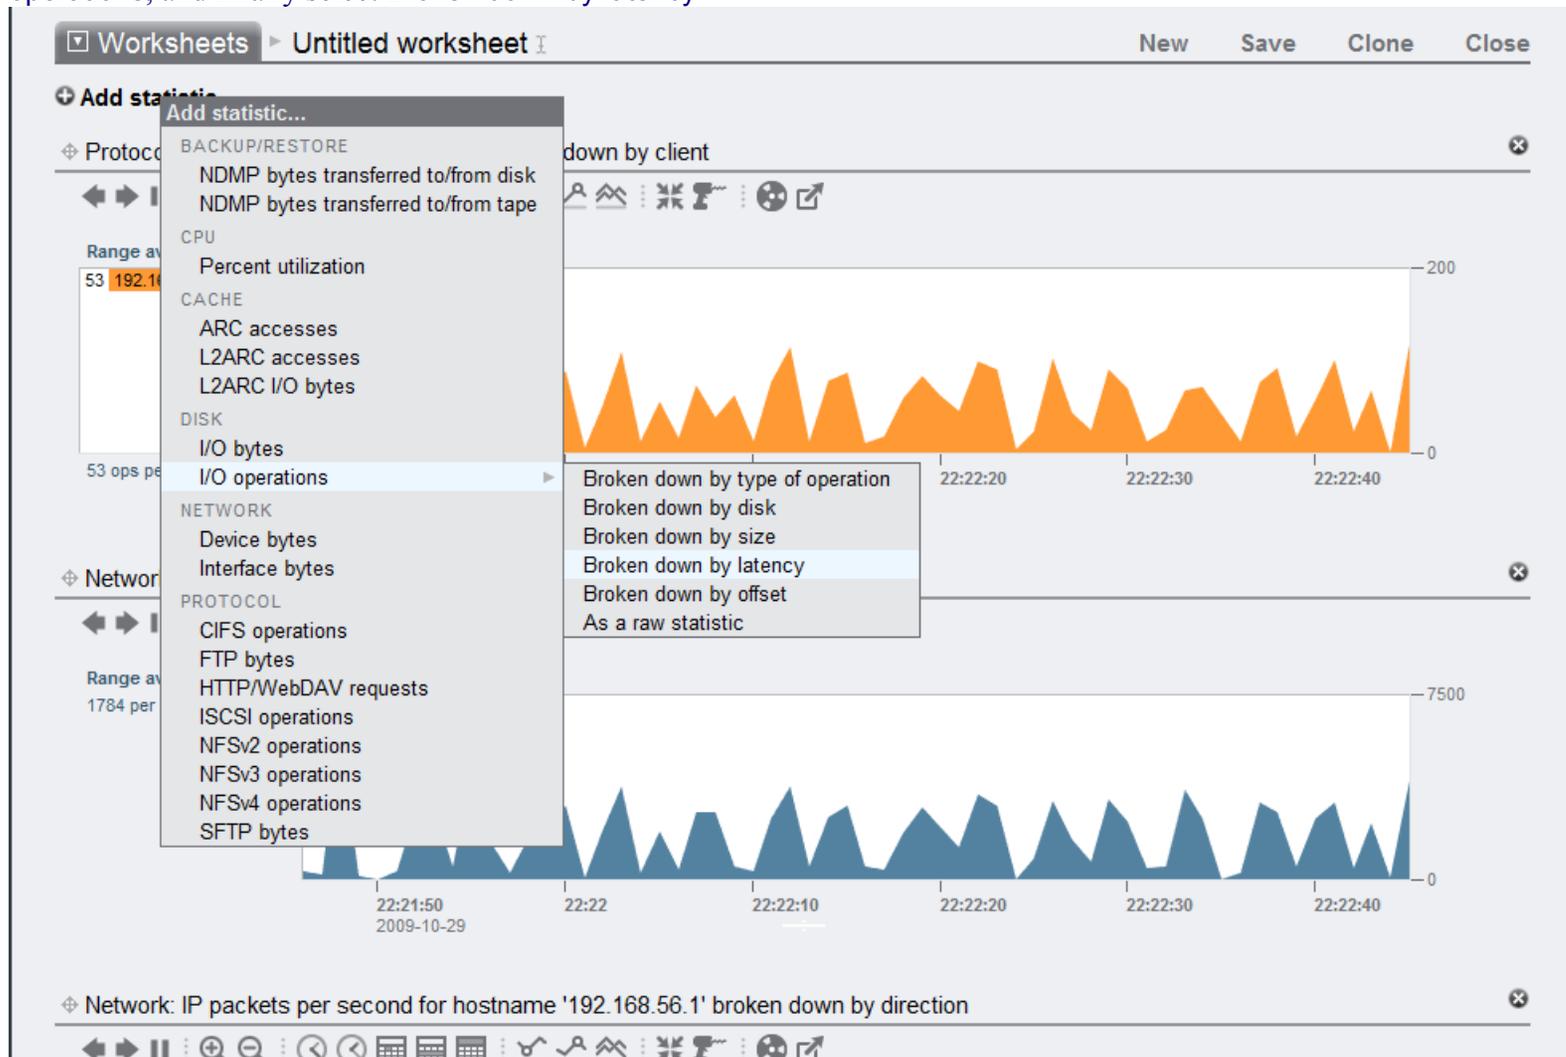
4. a second statistics is added to the worksheet showing traffic by IP packets



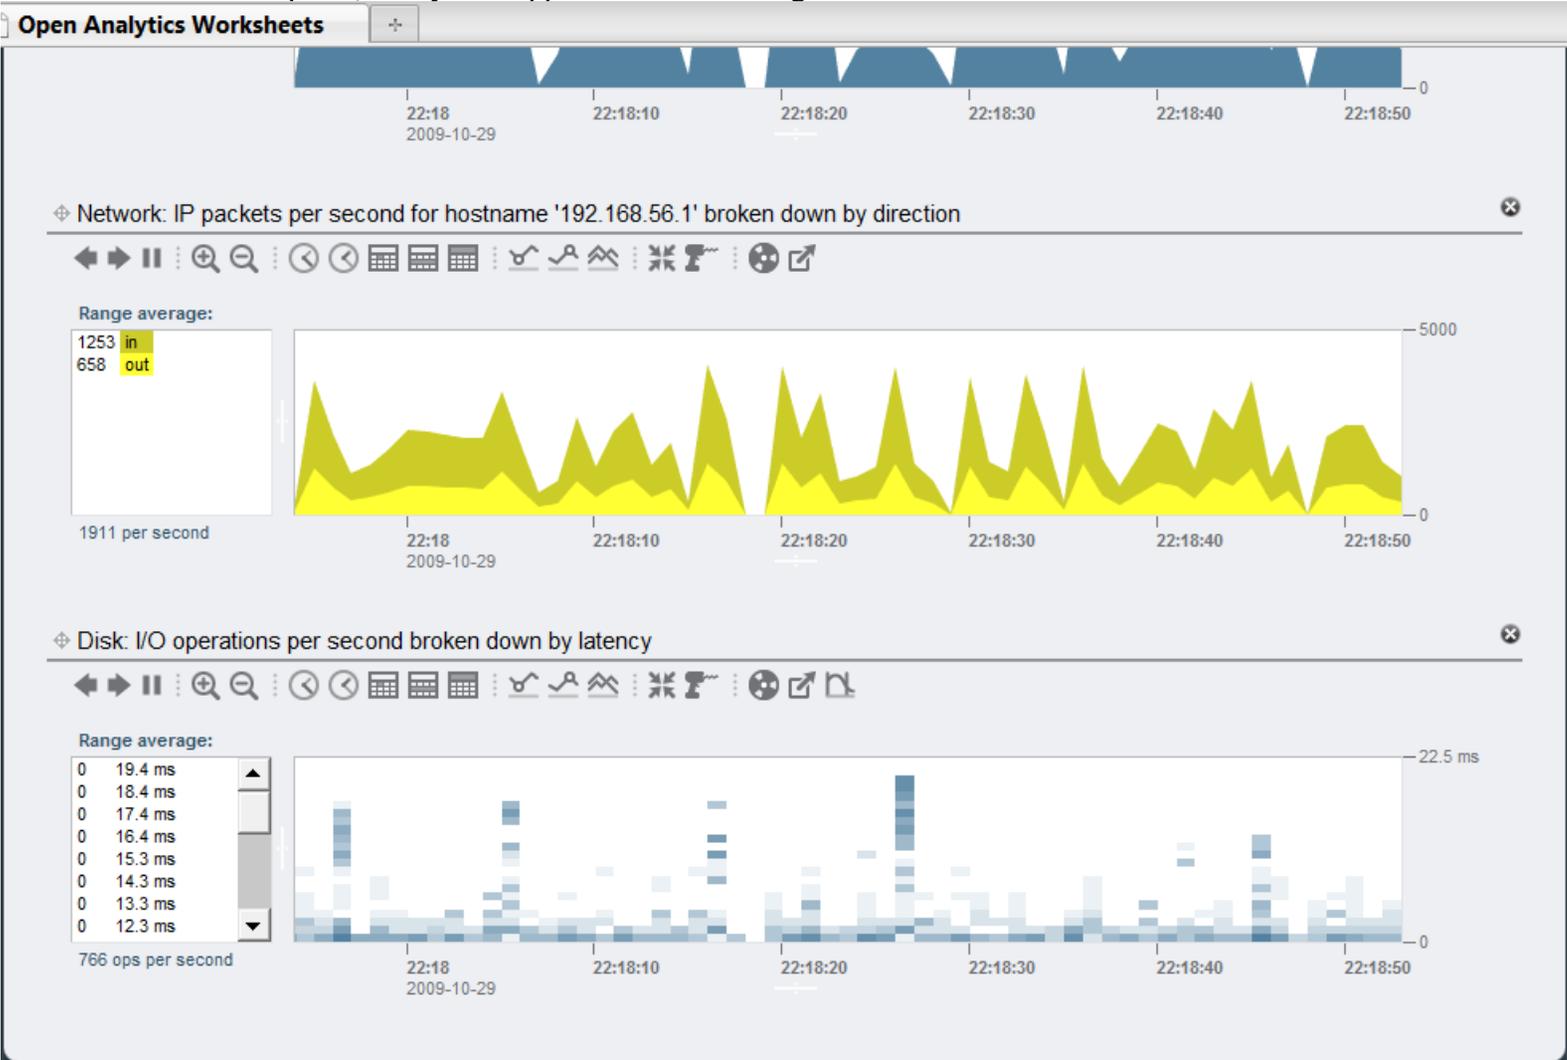
5. to demonstrate the drill down capabilities further, click on the power drill icon in the second statistic and select By direction
6. in the Range average list, multi-select both in (single click on the word in) and out (hold down the shift key and single click on the word out)



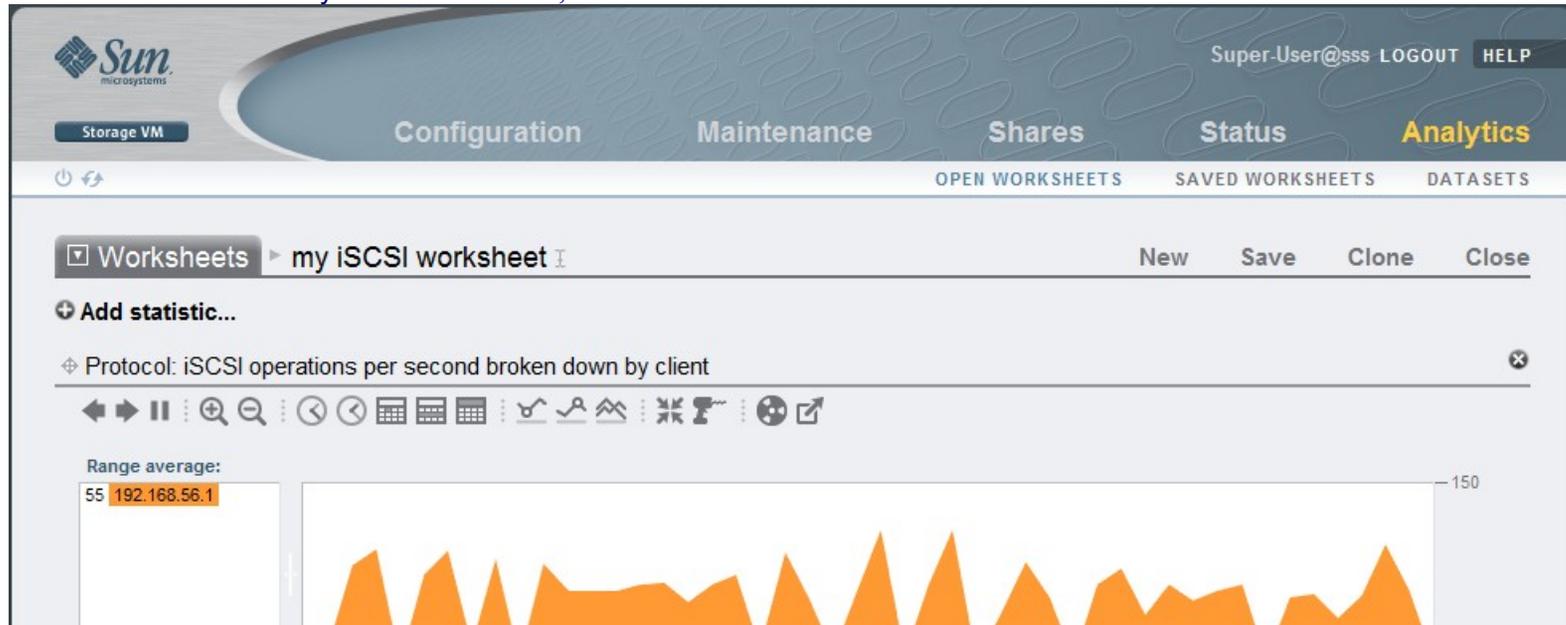
7. while the system is busy it may be handling the load without incident
8. to understand performance from the use perspective, go back to the top of the Worksheet, click on Add statistic ... , select I/O operations, and finally select Broken down by latency



- 9. you can now observe the effects of the load on performance from the users' perspective
- 10. while there are a few spikes, the system appears to be servicing the users well



11. to save this worksheet for future use, go to the top of the Analytics area and select the text **Untitled worksheet**, then enter a title for the worksheet such as **my iSCSI worksheet**, and click on the **Save** label



When you down looking at the statistics, you need to do some cleanup work for the next part of the lab

1. select **Close** on the **BUI** to close the worksheet (you saved it and can get it back later of you wish)
2. go back to the **Windows** desktop, select the command shell running the **start_iSCSI_load.bat** script, and press **<CTRL>-C** to kill the load generator

C. CIFS Attached Storage

1. Generate Load

First, we need to create some load on the storage system (files being written back and forth to the allocated storage).

1. go to the Windows desktop in your lab environment and open (double-click on) a folder named `lab_utils`
2. inside this folder should be 3 files



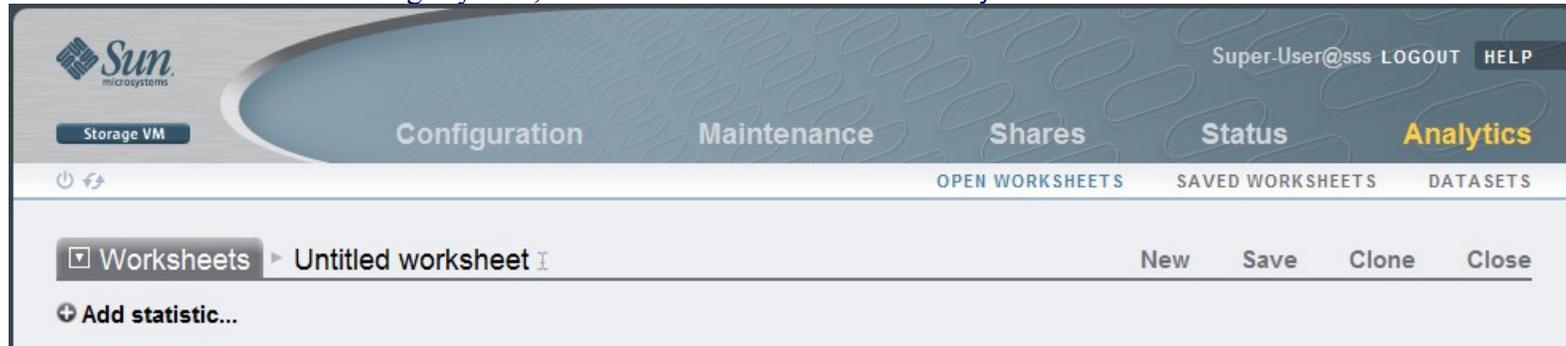
3. open or start (by double-clicking on it) the file named `start_CIFS_overload.bat`

Your Windows client is now trying to fill up the CIFS share on the storage system. This program takes about 2 minutes to completely fill up the drive space. At that point it will pause for about 5 seconds, then clear the CIFS shares of previous files, and then start filling the CIFS share up all over again.

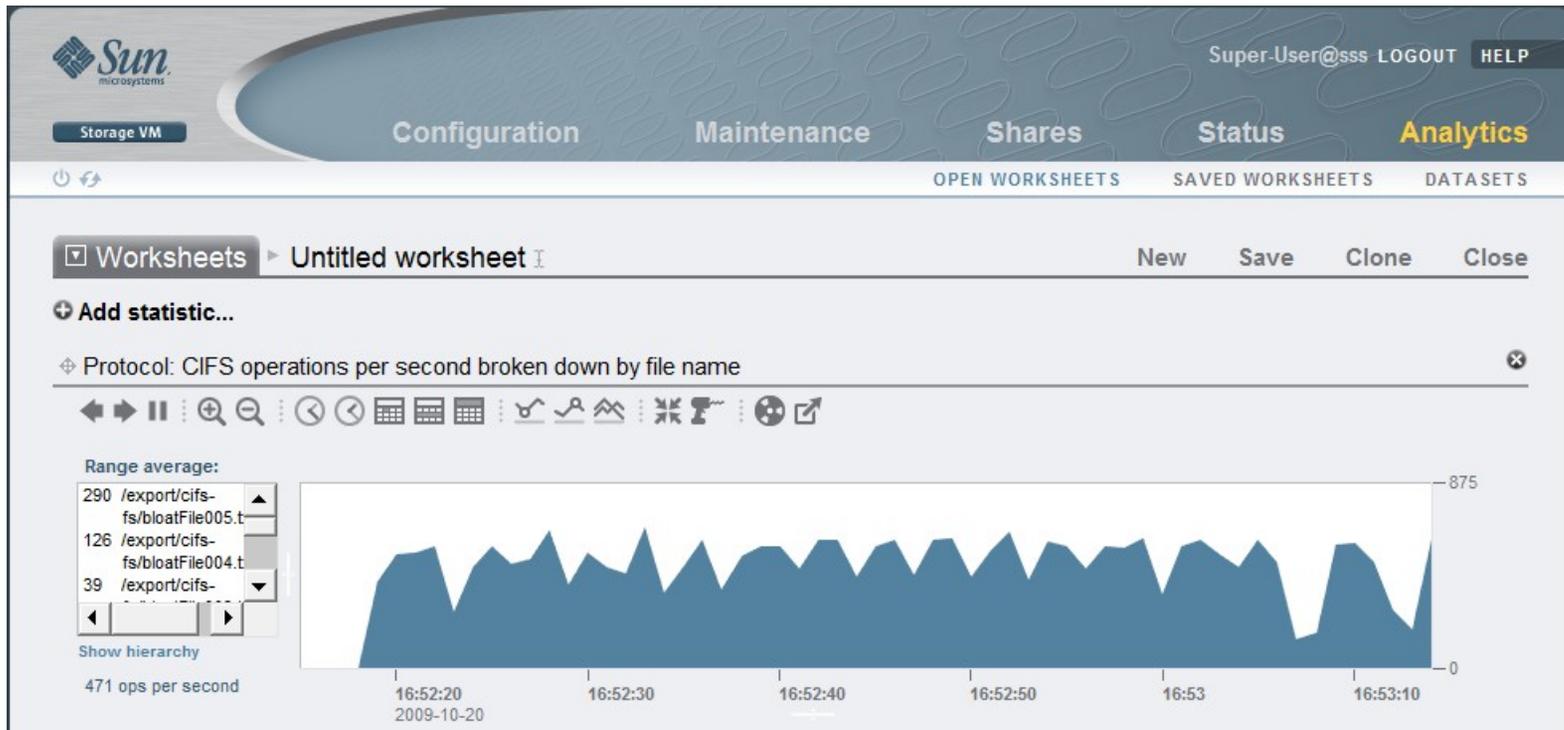
2. Setup Worksheets

Second, you will want to setup an Analytics worksheet to observe the storage system's behavior under the load being generated on it.

1. in the Firefox BUI for the storage system, click on the menu item titled Analytics



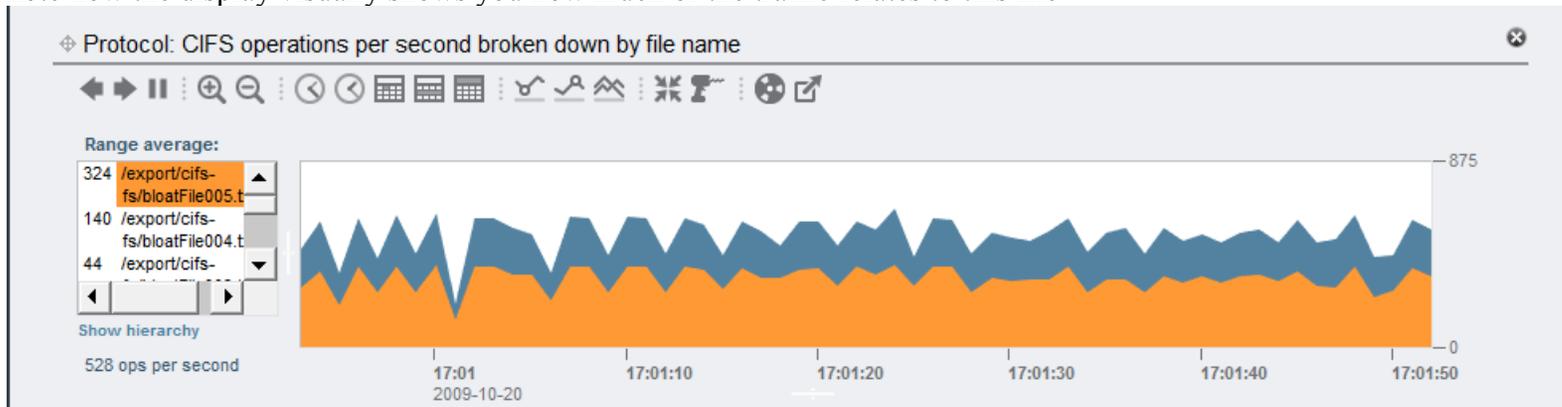
2. you get an empty untitled worksheet
3. start to add a Statistic by clicking on the Add statistic ... label
4. the popup menu shows you the high-level categories for available statistics
5. navigate in the popup menu to the item CIFS operations → Broken down by file name and select it



6. a lot of traffic hitting the CIFS share
7. in the Range average list you see file bloatFile005.txt taking up a lot of operations per second

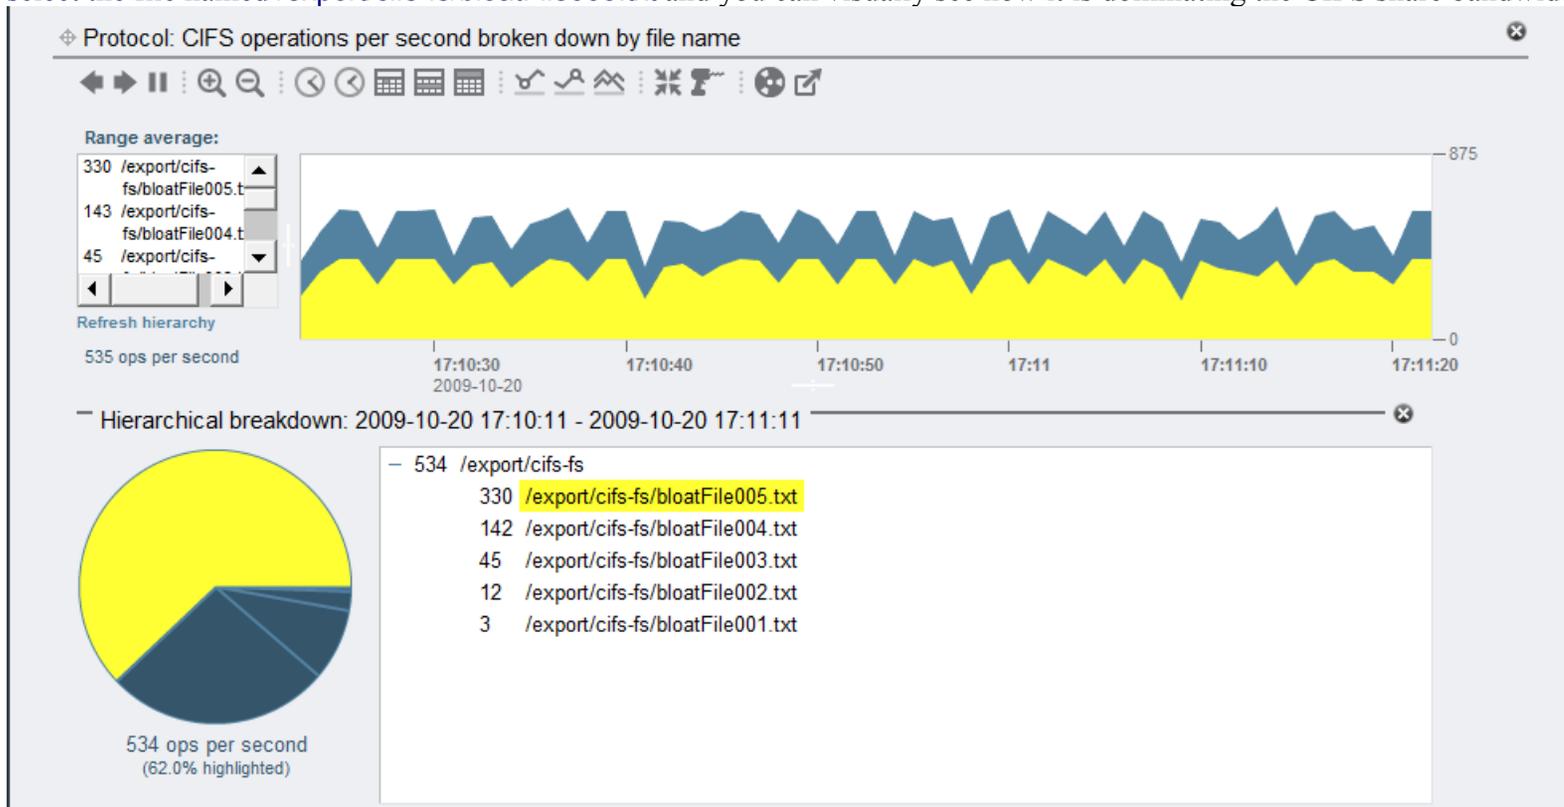
3. Drill Down

1. click on the file `bloatFile005.txt` in the Range average list
2. note how the display visually shows you how much of the traffic relates to this file



3. underneath the Range average list is an option to Show hierarchy; go ahead and select this option

4. new detail appears showing a pie chart, list of active files, and the operation per second per file
5. select the file named `/export/cifs-fs/bloatFile005.txt` and you can visually see how it is dominating the CIFS share bandwidth



This is a small sampling of the details you can examine on the storage systems operation and usage. Feel free to use the load generation scripts and other statistics to further explore the Analytics capabilities of the Sun Storage 7000 Unified Storage System.