

Installing a perfSONAR Network Performance Node

Rich Carlson

NWS Workshop

Jan 12-14, 2009



Introduction

- Most of the Internet2 network performance tools are client/server based
 - Client software runs on the users Laptop/Desktop computer
 - Software is automatically downloaded when a test is requested
 - Server software runs on a dedicated server which is physically attached to the network on the same LAN as the production server
 - Software is pre-installed and must be customized before using

Initial Suite of Tools

- NDT – a client/server program to find host configuration and network infrastructure faults
- NPAD – a client/server program to verify host configuration and local infrastructure will work with remote sites
- BWCTL – a client/server program to continuously monitor the network infrastructure
- OWAMP – a client/server program to continuously monitor the end-to-end delay

perfSONAR

- A framework for building large scale, multi-domain network measurement and monitoring infrastructures
 - Reference implementations of tools that use this framework have been developed by Internet2 engineers to demonstrate the value of this approach

perfSONAR perfNode

- A pre-built ISO image that can be booted to quickly bring up a measurement server
 - Initial boot requires configuration and customization steps (this presentation)
 - Subsequent boots bring the system back to this known state

- Most of the data in this slide deck comes from the on-line quick-start guide by Jason Zurawski <http://code.google.com/p/perfsonar-ps/wiki/NPTToolkitQuickStart>

Step 1 – Obtain the ISO image

- Download the latest ISO image from <http://www.perfsonar.net/download/bundles/pS-NPToolkit/current.iso>
 - OR, ask Internet2 staff to send you a CD-ROM disk
- Burn the downloaded ISO image onto a CD-ROM disk
 - Be sure and burn this as an ISO image and not as a data disk, see the instructions for your burner software.

Step 2 – Prepare the Hardware

- Find a computer you can dedicate as a server
 - 1.8 GHz Intel/AMD processor
 - 512 MB RAM
 - 80 GB Hard disk (Disk needs to be formatted)
 - 100 Mbps Ethernet interface
 - CD read-only disk drive

- Note: a low-cost rack-mount server is available from KOI computers <http://www.koicomputers.com/> for approximately \$600, contact Internet2 staff for details

Step 3 – Initial Boot

- Insert the CD-ROM disk into the drive and power on the system. The following screen will appear. Press “Enter” or wait 30 seconds for boot process to begin

```
ISOLINUX 3.11 2005-09-02 Copyright (C) 1994-2005 H. Peter Anvin
Press keys F2 or F3 for help and boot options.          (Based on Knoppix V5.
0.1-CD)

Performance Node: LiveCD
Integrated by Internet2

http://www.internet2.edu/performance/

boot: _
```

Step 4 – Begin Configuration

- After a short time, the following screen will appear. All items in RED need to be configured.

```
Internet2 Network Performance Toolkit customization script
Tools in RED need to be customized
Tools in MAGENTA are not required

1. Configure drive to hold data/customizations
2. Configure site settings
3. Set built-in account passwords
4. Enable/Disable Services
5. Manage Users
6. Configure BWCTL
7. Configure NTP
8. Configure OWAMP
9. Configure Static IP
10. Change Timezone
11. Modify advertised interface
0. exit

Make a selection:
```

Step 5 – Select a Drive

- Enter '1' to configure the Drive

```
Make a selection: 1

Welcome to the Internet2 Performance Node LiveCD drive configuration program.
This program will prompt you to setup your harddrives/removable media so they can be used with the toolkit.
Note: You need to format/partition your drives before running this script.
Default values are in []

Please choose a location for where you would like to save the toolkit's configuration:
  1. /dev/sda      Internal Hard Drive      Capacity: 2.2G      83% free
  0. exit

Please select a drive: _
```

- In this case only 1 drive is available, select the drive listed on your server
 - Note: the drive will be reformatted if necessary

Step 6 – Configure Site Settings

- Select '2' to enter the Site Settings option
 - Enter admin's name/email, org name and physical location, and the link speed, use "Internet2 Arts_and_Humanities" as the communities of interest

```
Make a selection: 2
```

```
Welcome to the Internet2 Performance Node LiveCD site information configuration program.  
This program will prompt you for your site information.  
Default values are in []
```

```
Enter your administrator's name [Joe Admin]: Jason Zurawski
```

```
Enter your site name [Some Organization]: Internet2
```

```
Enter your site's location [Some City, State, Country]: Washington, DC, USA
```

```
Enter the communities of interest, separated by spaces (e.g. LHC Internet2 eULBI).
```

```
If you're confused as to what to enter, visit http://code.google.com/p/perfsonar-ps/wiki/NPToolkit
```

```
[]): LHC Internet2 MAX
```

```
Please choose your connection speed:
```

```
1. 100 Mbps
```

```
2. 1 Gbps
```

```
3. 10 Gbps
```

```
Select your connection speed [2]: 1
```

```
Enter an email to be used for trouble reporting [joe.admin@organization.edu]: zurawski@internet2.edu
```

Step 7 – Enter passwords

- Enter '3' to set the system passwords, you will not be able to use the system if you don't set these passwords

```
Make a selection: 3
```

```
Welcome to the Internet2 Performance Mode LiveCD Password configuration program.  
This program will help you set the passwords for built-in accounts: knoppix and root  
Default values are in []
```

```
The password for user knoppix needs to set.  
Please specify a password for the knoppix account.
```

```
Enter new UNIX password:
```

```
Retype new UNIX password:
```

```
passwd: password updated successfully
```

```
The password for user root needs to set.
```

```
Please specify a password for the root account.
```

```
Enter new UNIX password:
```

```
Retype new UNIX password:
```

```
passwd: password updated successfully
```

Step 8 – Configure specific services

- Enter '4' to enable/disable specific services
 - Enable NDT, NPAD, BWCTL, OWAMP and Web Service
 - Enable SSH only if you want to remotely login to this server
 - Disable PingER, SNMP, and perfSONAR-BUOY services

```
Make a selection: 4

Welcome to the Internet2 Performance Node LiveCD service configuration program.
This program will let you select which services you want to enable.
You may enable or disable any service or groups of services
 1 Disable Stock Bandwidth Services (BWCTL, NDT, NPAD)
 2 Disable Stock Latency Services (OWAMP, PingER)
 3 Enable SSH
 4 Disable Web Services
 5 Disable BWCTL
 6 Disable OWAMP
 7 Disable NDT
 8 Disable NPAD
 9 Disable PingER
10 Disable SNMP MA
11 Disable perfSONAR-BUOY
12 Enable All Services
 0. exit

Make a selection: _
```

Step 9 – Manage Users

- Skip this item, unless you specifically need to add a new user, you can skip this item!

Step 10 – Configure BWCTL

- Enter '6' and press 'Enter' when prompted. Simply take the default settings

```
Make a selection: 6

Welcome to the Internet2 Performance Mode LiveCD BWCTL configuration program.
This program will help you customize BWCTL to meet your needs.
Default values are in []

Please provide some basic information about your system...
What is the maximum amount of bandwidth (in Mbps) that should be usable for testing? (to omit type *) [900]: *

The next few questions let you specify what tests you are willing to let users perform
What is the maximum test duration (in seconds)? (to omit type *) [60]: *
What is the maximum number of pending requests to queue? (defaults to unlimited):
How many seconds should pending requests be allowed to queue? [1800]:

Advanced options...
You may grant specific networks unrestricted access now.
Enter a network and mask [ex. 192.168.0.0/16] (blank line to stop):

You may forbid specific networks from performing tests now.
Enter a network and mask [ex. 192.168.0.0/16] (blank line to stop):

You may now grant specific users unrestricted access. You will need to provide a username & password for each user.
To delete a user, select their username from the list below.
    1. Add a new user
    0. Exit
Make a selection: _
```

Step 11 – Configure NTP

- Enter '7', then simply enter '0' to exit this page. Contact your campus IT staff if you want to change these servers

```
Make a selection: 7

Welcome to the Internet2 Performance Node LiveCD NTP configuration program.
This program will help you create a customized NTP configuration for your site.
Default values are in []

You can select the NTP servers from the below list. Currently selected NTP servers have a '*' next to them.
You should have 5 or 6 NTP servers selected to maximize NTP's effectiveness.
You can specify all the NTP servers to use by entering the options separated by spaces (e.g. 2 3 4 5 6)

1. Custom NTP Server
2. Internet2 - Atlanta, GA USA
3. Internet2 - Chicago, IL USA *
4. Internet2 - Houston, TX USA *
5. Internet2 - Salt Lake City, UT USA
6. Internet2 - Los Angeles, CA USA *
7. Internet2 - New York, NY USA *
8. Internet2 - Sunnyvale, CA USA *
9. ESnet - New York, NY USA *
10. ESnet - Sunnyvale, CA USA *
11. NIST - Gaithersburg, MD USA
12. Naval Observatory - Washington, DC USA
13. Naval Observatory - Columbus, OH USA
14. Naval Observatory - Colorado Springs, CO USA
15. Naval Observatory - Los Angeles, CA USA
16. Naval Observatory - Seattle, WA USA
17. Naval Observatory - Fairbanks, AK USA
18. Maui HPC Center - Maui, HI USA
0. exit

Make a selection: 0
```

Step 12 – Configure OWAMP

- Enter '8' and then press 'Enter' when prompted to take the default values

```
Make a selection: 8

Welcome to the Internet2 Performance Node LiveCD OWAMP configuration program.
This program will help you customize OWAMP to meet your needs.
Default values are in []

Please provide some basic information about your system...
What is the maximum amount of bandwidth (in Mbps) that should be usable for testing? (to omit type *) [1]: *
What is the maximum amount of disk space (in megabytes) that should be usable for testing? (to omit type *) [500]: *

The next few questions let you specify what tests you are willing to let users perform
Should buffered data files be deleted automatically as they are fetched? [yes]:

Advanced options...
You may grant specific networks unrestricted access now.
Enter a network and mask [ex. 192.168.0.0/16] (blank line to stop):

You may forbid specific networks from performing tests now.
Enter a network and mask [ex. 192.168.0.0/16] (blank line to stop):

You may now grant specific users unrestricted access. You will need to provide a username & password for each user.
To delete a user, select their username from the list below.
    1. Add a new user
    0. Exit
Make a selection:
```

Step 13 – Configure Static IP address

- Only take this step if your campus IT staff provides you with a static IP address. In most cases your campus will automatically assign an IP address to the server during the boot process

```
Make a selection: 9
```

```
Welcome to the Internet2 Performance Mode LiveCD Static IP configuration program.  
This program will prompt you to enable/disable a static IP address.  
Default values are in []
```

```
After you have configured the static IP, it can be disabled by selecting "Configure Static IP" again. You will then be prompted to disable it.
```

```
Enter the network interface identifier (Enter * to cancel) [eth0]:
```

```
Enter the IP you would like [192.168.1.99]:
```

```
Enter your netmask [255.255.255.0]:
```

```
Enter your gateway's IP address [192.168.1.1]:
```

```
Enter your nameserver's IP address (3 max, blank line to exit early. If you enter none, the existing nameserver configuration will be used):
```

```
Do you want to start this interface (this may take a little while)? [Yes]:
```

Step 14 – Enter timezone info

- An optional step to set the timezone for your physical location.

```
Make a selection: 10
Current timezone: EST/EDT
Your current time zone is set to Etc/UTC
Do you want to change that? [n]: y

Please enter the number of the geographic area in which you live:

    1) Africa                7) Australia
    2) America              8) Europe
    3) US time zones        9) Indian Ocean
    4) Canada time zones    10) Pacific Ocean
    5) Asia                  11) Use System V style time zones
    6) Atlantic Ocean       12) None of the above

Then you will be shown a list of cities which represent the time zone
in which they are located. You should choose a city in your time zone.

Number: _
```

Step 15 – Modify advertised interface

- Skip this item. In most cases your server will have only 1 network interface so you can skip this task.

Configuration Complete

- Once you have configured all these items, Press '0', the system may reboot, or it may just continue on and complete the boot process.
- When the boot process is complete you will receive a login prompt.
 - The login id is "knoppix" (no quotes)
 - The password is what you assigned

User Login Complete

```
Creating SSH2 RSA key; this may take some time ...
Creating SSH2 DSA key; this may take some time ...
Starting OpenBSD Secure Shell server: sshd.
Starting bwctld.
bwctld[3724]: FILE=time.c, LINE=123, NTP: Status UNSYNC (clock offset problems likely)
owampd started
Starting NDT Services:      fakewww web100srv.
Starting NPAD serverStarting owampd.
owampd[3747]: FILE=time.c, LINE=112, NTP: Status UNSYNC (clock offset issues likely)
owampd started
/usr/local/perfSONAR-PS/perfsonar-daemon.pl --config /usr/local/etc/perfSONAR/PingER.conf --logger=/usr/local/etc/perfSONAR/Pi
ER_logger.conf --pidfile=PingER.pid --logger=/usr/local/etc/perfSONAR/PingER_logger.conf --user=perfsonar --group=perfsonar 2>
dev/null
/etc/init.d/PingER.sh start: PingER started
/usr/local/perfSONAR-PS/perfsonar-daemon.pl --config /usr/local/etc/perfSONAR/pSB_MA.conf --logger=/usr/local/etc/perfSONAR/pSB
MA_logger.conf --piddir=/var/run --pidfile=pSB_MA.pid --user=perfsonar --group=perfsonar
/etc/init.d/pSB.sh start: pSB started
/usr/local/perfSONAR-PS/perfsonar-daemon.pl --config /usr/local/etc/perfSONAR/SNMP_MA.conf --logger=/usr/local/etc/perfSONAR/S
NMP_MA_logger.conf --piddir=/var/run --pidfile=SNMP_MA.pid --user=perfsonar --group=perfsonar
/etc/init.d/snmpMA.sh start: SNMP MA started
Conf::must_get_sublist(LIST=>NODE, ) undefined, called from /usr/local/AMI/script/bwcollector.pl:238
Couldn't start perfSONAR-BUOY Measurement Collector Daemon
Conf::must_get_val(ATTR=>NODE, ) undefined, called from /usr/local/AMI/script/bwmaster.pl:154
Couldn't start perfSONAR-BUOY Collection Master
Starting anac(h)ronistic cron: anacron.
Starting periodic command scheduler: crond.
Starting web server (apache2)...apache2: Could not reliably determine the server's fully qualified domain name, using 127.0.0.
for ServerName
.
Setting Local options: Turning off TSO optionCannot get device rx csum settings: Operation not supported
Cannot get device udp large send offload settings: Operation not supported
.
Debian GNU/Linux 4.0 Knoppix tty1

Knoppix login: knoppix
Password:

Welcome to the Performance Node Live CD v2.0!!

To perform site re-configuration run 'sudo nptoolkit-configure.py' .

knoppix@tty1[~]1$
```

Verify IP address

- Enter the command 'ifconfig' to verify IP address is configured and operational

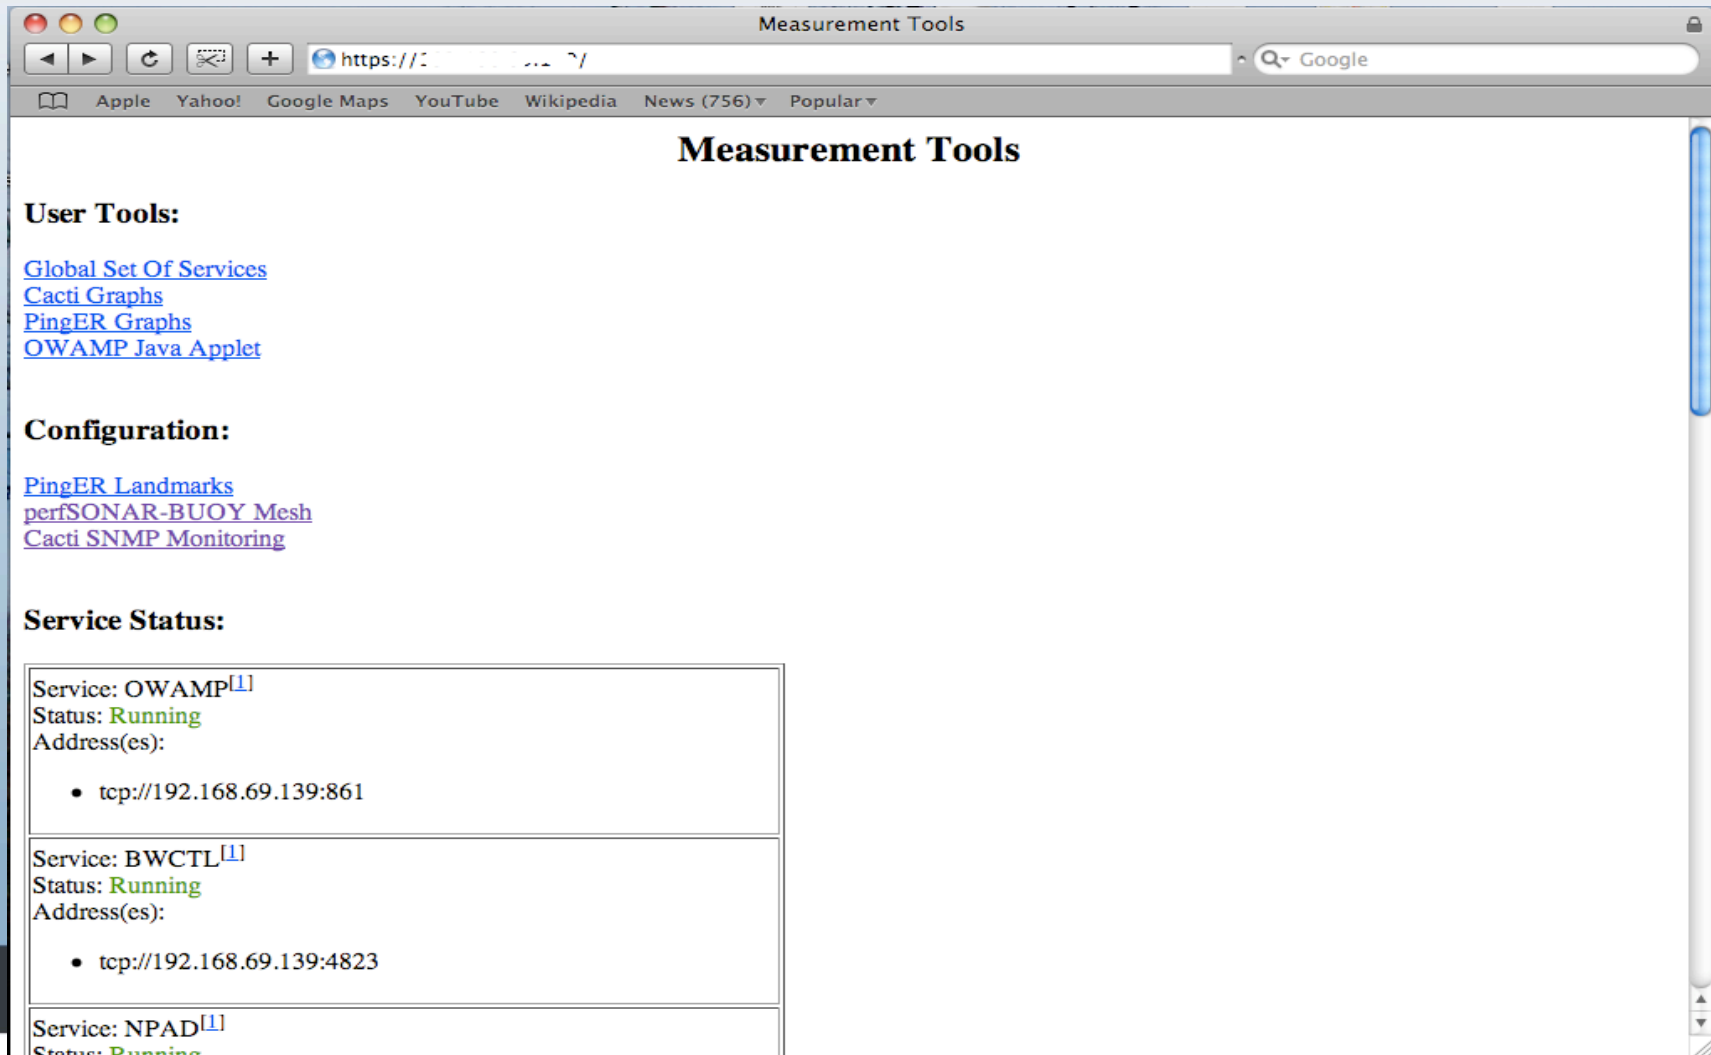
```
zurawski@tty1[~]$: ifconfig
eth0      Link encap:Ethernet  HWaddr 00:0C:29:F3:89:0E
          inet addr:192.168.69.139  Bcast:192.168.69.255  Mask:255.255.255.0
          inet6 addr: fe80::20c:29ff:fef3:890e/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:1213 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1287 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:543741 (530.9 KiB)  TX bytes:545482 (532.6 KiB)
          Interrupt:16 Base address:0x2000

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:606 errors:0 dropped:0 overruns:0 frame:0
          TX packets:606 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:99822 (97.4 KiB)  TX bytes:99822 (97.4 KiB)
```

Verify Operation

- From another host bring up a web browser and enter <https://{ip-address}> where {ip-address} is the address assigned to this server in the w.x.y.z format
- If your IT department has assigned a DNS name to this address you can use that name instead.

Main Server Web page



The screenshot shows a web browser window with the title "Measurement Tools". The address bar contains a URL starting with "https://". The browser's search bar shows "Google". The page content is organized into three main sections:

- Measurement Tools** (Section Header)
- User Tools:**
 - [Global Set Of Services](#)
 - [Cacti Graphs](#)
 - [PingER Graphs](#)
 - [OWAMP Java Applet](#)
- Configuration:**
 - [PingER Landmarks](#)
 - [perfSONAR-BUOY Mesh](#)
 - [Cacti SNMP Monitoring](#)
- Service Status:**

Service: OWAMP ^[1] Status: Running Address(es): <ul style="list-style-type: none">• tcp://192.168.69.139:861
Service: BWCTL ^[1] Status: Running Address(es): <ul style="list-style-type: none">• tcp://192.168.69.139:4823
Service: NPAD ^[1] Status: Running

Finding other Server

- The “Global Set Of Services” link gets you to a new page that lists your peer servers.
- All remote servers should be listed, servers run by peer Arts & Humanities sites will appear under the “Arts_and_Humanities” heading

perfSONAR Global Service and Data View

Table Of Contents

[Measurement Tools](#)

[NPAD Daemons](#)

[NDT Daemons](#)

[OWAMP Daemons](#)

[BWCTL Daemons](#)

[perfSONAR Services](#)

[SNMP MA \(Utilization\)](#)

[perfSONAR-BUOY \(Iperf Data\)](#)

[perfSONAR-BUOY \(OWAMP Data\)](#)

[PingER](#)

NPAD Servers

ESnet NPAD Server	http://198.128.1.239:8200	NPAD Server at ESnet in Berkeley, CA
ESnet NPAD Server	http://198.128.1.239:8200	NPAD Server at ESnet in Berkeley, CA
Internet2 NPAD Server	http://207.75.164.246:8200	NPAD Server at Internet2 in Ann Arbor, MI USA

NDT Servers



Server Status

- Your server is now ready to support testing
- You can point clients to it so verify your site is operating properly
- You can point clients to the remote site to verify that the WAN path is capable of carrying your traffic

Running NDT test

- Press “Start” Button

ESnet NDT server

http://:7123/

Apple Yahoo! Google Maps YouTube Wikipedia News (3294) Popular

perfAdmin MA Output Cacti Traffic Analy... DAMSL CIS ... Cacti NPToolkitQu... kerberos - ... ESnet NDT s...

ESnet Web100 based Network Diagnostic Tool (NDT)

Located at Ames, Iowa, USA; 1000 Mbps (Gigabit Ethernet) network connection

This java applet was developed to test the reliability and operational status of your desktop computer and network connection. It does this by sending data between your computer and this remote NDT server. These tests will determine:

- The slowest link in the end-to-end path (Dial-up modem to 10 Gbps Ethernet/OC-192)
- The Ethernet duplex setting (full or half);
- If congestion is limiting end-to-end throughput.

It can also identify 2 serious error conditions:

- Duplex Mismatch
- Excessive packet loss due to faulty cables.

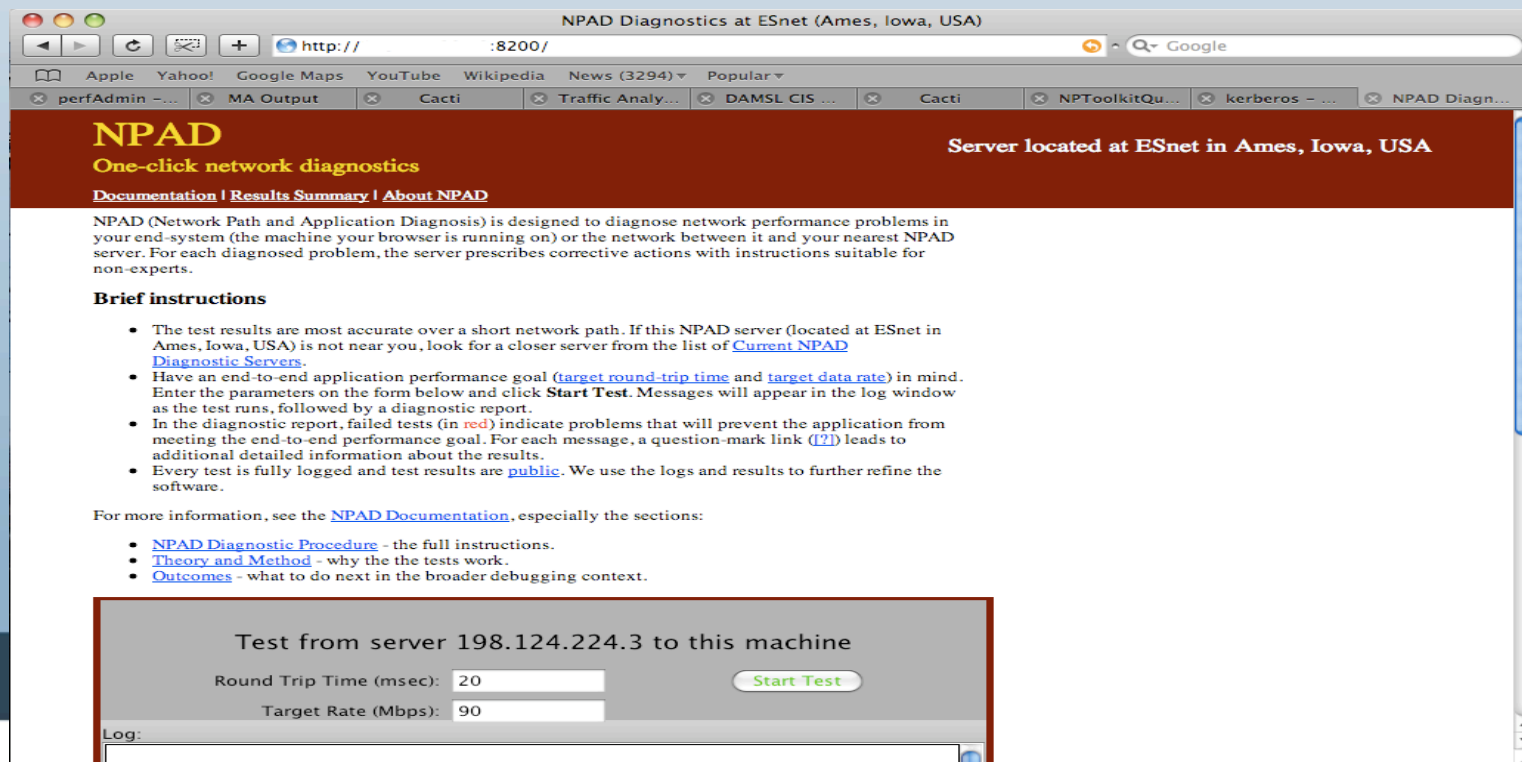
A test takes about 20 seconds. Click on "start" to begin.

TCP/Web100 Network Diagnostic Tool v5.5.4a
click START to begin

START Statistics More Details... Report Problem Options

Running NPAD test

- Test to Local server and enter required speed and RTT to remote location then press “start”



The screenshot shows a web browser window titled "NPAD Diagnostics at ESnet (Ames, Iowa, USA)". The address bar shows "http://:8200/". The browser tabs include "perfAdmin", "MA Output", "Cacti", "Traffic Analy...", "DAMSL CIS...", "Cacti", "NPToolkitQu...", "kerberos", and "NPAD Diagn...". The page content includes the NPAD logo, a header "Server located at ESnet in Ames, Iowa, USA", and a "Brief instructions" section. A form at the bottom is highlighted with a red box, showing "Test from server 198.124.224.3 to this machine" with input fields for "Round Trip Time (msec): 20" and "Target Rate (Mbps): 90", and a "Start Test" button. A "Log:" field is visible below the form. The ESnet logo is in the bottom right corner.

NPAD
One-click network diagnostics
Server located at ESnet in Ames, Iowa, USA
Documentation | Results Summary | About NPAD

NPAD (Network Path and Application Diagnosis) is designed to diagnose network performance problems in your end-system (the machine your browser is running on) or the network between it and your nearest NPAD server. For each diagnosed problem, the server prescribes corrective actions with instructions suitable for non-experts.

Brief instructions

- The test results are most accurate over a short network path. If this NPAD server (located at ESnet in Ames, Iowa, USA) is not near you, look for a closer server from the list of [Current NPAD Diagnostic Servers](#).
- Have an end-to-end application performance goal ([target round-trip time](#) and [target data rate](#)) in mind. Enter the parameters on the form below and click **Start Test**. Messages will appear in the log window as the test runs, followed by a diagnostic report.
- In the diagnostic report, failed tests (in red) indicate problems that will prevent the application from meeting the end-to-end performance goal. For each message, a question-mark link (?) leads to additional detailed information about the results.
- Every test is fully logged and test results are [public](#). We use the logs and results to further refine the software.

For more information, see the [NPAD Documentation](#), especially the sections:

- [NPAD Diagnostic Procedure](#) - the full instructions.
- [Theory and Method](#) - why the the tests work.
- [Outcomes](#) - what to do next in the broader debugging context.

Test from server 198.124.224.3 to this machine

Round Trip Time (msec):

Target Rate (Mbps):

Log:

Conclusion

- Server should be permanently installed at your home institution 'on' your local network
- Once the ISO image is booted, the tools will register with a global lookup service so others can find them
- These tools will make it easier to find and fix network performance problems