

CS 638 Lab 1: Introduction to Schooner

Joe Chabarek and Paul Barford

University of Wisconsin –Madison

jpchaba,pb@cs.wisc.edu

The Wisconsin Advanced Internet Laboratory (WAIL) is a one-of-its-kind facility for conducting network research. Through the support of the TOSA Foundation, Cisco Systems, the National Science Foundation and other partners, WAIL is now widely used by research groups all over the world. Throughout the semester, you will be using WAIL to conduct experiments that give you access and hands-on experience with the wide variety of networking equipment that is available in WAIL. The web-based portal that gives you remote, real-time access to WAIL and all of its capabilities is called *Schooner*. This first lab is intended to get you up and running on creating topologies in Schooner and to give you hands-on experience with simple point-to-point networks.

1 Overview and Objectives

Lab 1 is divided into two parts. The first part will introduce you to Schooner. The second part will require you to go to WAIL so you can build an Ethernet cable, and physically configure a home local area network.

Upon completion of the lab exercises below, students will be able to:

1. *Create simple experimental topologies via the Schooner portal in the WAIL laboratory*
2. *Save and recall an experimental configuration, and to access host nodes in an instantiated experiment.*
3. *Conduct simple communications tests between*

nodes in an experiment to ensure they are connected.

4. *Make a straight-through Ethernet cable.*
5. *Create a simple Local Area Network using a hub, cables and Windows OS such as might be used in a home environment.*

2 Part 1: Introduction to Schooner

This part of lab 1 will take you through the tasks of creating an account in Schooner, creating an experiment, saving and recalling an experiment and conducting tests in an experiment. While the basic requirements of the lab are relatively procedural, students are encouraged to create additional experiments and to consider what else might be necessary to expand beyond simple topologies in order to create experiments that have larger scale and diversity.

2.1 Step 1: Creating an Account

Access the Schooner homepage at <http://www.schooner.wail.wisc.edu>. Follow the instructions on the Request Account page since you will need the account in order to access the WAIL lab equipment. Processing the account may take up to 24 hours, since it entails a manual approval process by the WAIL lab managers. The purpose for manual approval is to ensure that the infrastructure is used for research by scientists and students under the supervision of their professors.

There are a number of resources that will provide you with introductory and tutorial information on the lab and the Schooner interface. For example, please access the Schooner FAQ, How to Get Started, and the Emulab FAQ, from the 'Documentation' link on the Schooner Web page for more details. From a broader perspective, you are encouraged to investigate all of the links to learn more about the infrastructure and the projects that are currently underway.

At the completion of the account request process, you will be able to access Schooner, create experiments and then access the equipment used in experiments. You will receive email once your account has been approved. You can then log into your account by going to the Schooner page. Click the Log In button on the left hand side. Initially you are shown your "My Emulab" page, which lists information about your account, including ongoing experiments.

2.2 Behind the scenes

A key design requirements for WAIL include the following:

1. The infrastructure support simultaneous use by multiple experiments (although the systems in use in any experiment will be exclusive to that experiment).
2. Users have the ability to remotely create arbitrary interconnection topologies that require no physical changes in the infrastructure (e.g., moving cables) in order to connect any two components.
3. Users have the ability to include actual networking equipment (e.g., routers) in their experiments.
4. Users have complete control over end hosts instantiated in their experiments - including the ability to reconfigure operating system images.
5. Users have the ability to emulate wide area link characteristics and generate representative traffic in their experiments.
6. Users have the ability to deploy high performance monitoring (e.g., Endace DAG packet monitors) in their experiments.

To achieve these requirements, a large data and configuration management infrastructure has been developed. This system manages user accounts, allocates available equipment and configures layer 1 - 3 networking components to enable various topologies of equipment. Beyond the physical and network connections, it is up to the users to ensure that any additional software components and configurations are properly set.

2.3 An Emulab-based Environment

Schooner is an extension of a tool called *Emulab* (<http://www.emulab.net>) that was designed to enable PCs to be connected together in different configurations. Once the PC topologies are established, direct SSH connections are used to access and configure them. The big difference between WAIL and Emulab is that all experiments in Emulab are *emulations of networks using PCs*, while WAIL has the ability to include actual networking hardware that runs real networking protocols. WAIL also has the ability to operate in an entirely emulated fashion, which is what will be done in lab 1.

2.4 Step 2: Start a New Experiment.

Once you have received your login information, log into Schooner (lower left since of the Schooner homepage). Next, click on the "Begin an Experiment" link in the lower left frame.

The default form allows you to upload an NS (Network Simulator) file. The NS format allows you to create complex network configurations through command line instructions. The NS format is actually a legacy from the NS2 network simulator which has been widely used over the years (see <http://www.isi.edu/nsnam/ns>). Luckily Schooner also includes a graphical topology construction tool that simplifies the topology generation process. Under the bullet "If you do not have a NS file" there is a link "New GUI editor". Click this link.

Use the GUI to create the physical setup/topology for the network shown in Figure 1. The interface is very simple to use. Simply drag an object into the workspace to create it. The example network shown in the figure has one LAN with two end hosts. To create this topology, please do the following:

1. Drag and drop the "Switch" icon onto the workspace. Notice the LAN Properties options on the right side of the screen. Leave the default values.
2. Drag and drop the "Node" icon onto the workspace. Again, their are properties to the right. Enter FC6-STD in the "OS" box (this results in the FC6 which is a Fedora Core 6 image to be loaded on your node). Leave the hardware box set to auto. Repeat this step to get two nodes in your network.
3. To establish links between the nodes, select the node icon in your workspace by clicking on it once. Start to move your mouse away from the node and it will start to draw a line, now click and hold the mouse and drag the line to the "lan0" icon. Connect your other host to the LAN similarly. Note that there are properties associated with the wire, including link speed. Feel free to play with these.
4. Click any item and then press the delete key to remove them.

2.5 A 2-host, LAN Network Configuration

Once the network is configured as directed, pull down the file menu and select "Create New Experiment". Give your experiment a name and a description, and click the Advanced tab to open it up. Change the default idle swap out time from four hours to two hours. This will cause your experiment to "swap out" after two hours, enabling resources to be more readily available to everyone. See the FAQ page for details on swapping. Finish by clicking "Create". By creating an experiment, you are actually generating an NS file that describes your configuration. You can also view

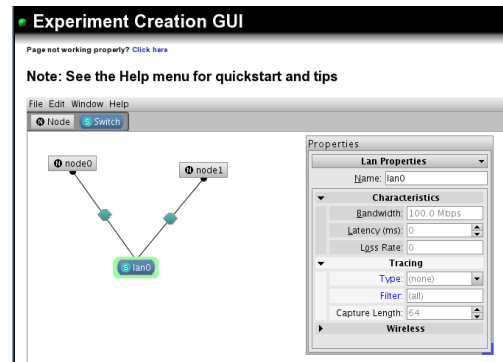


Figure 1: Network configuration/topology for part 1 shown in the NetBuild GUI.

the NS file in the GUI by pulling down the Window menu and selecting "NS File". See if you can identify how the file specifies your topology and the details of the nodes in the topology. This topology can now be recalled, modified and instantiated in WAIL and used for experiments.

2.6 Step 3: Instantiating Your Configuration

After creating the experiment from the GUI interface it will drop you into the experiment page. In the upper left, click "Swap Experiment In" to instantiate your experiment, it will take some number of minutes for the network to be setup. What is happening behind the scenes is that nodes are being reserved for you, your disk images are being loaded and your LAN topology is being configured on WAIL switches. You should be able to watch the control messages during the loading process (although most of the messages won't make much sense to you). As soon as the instantiation is complete you will receive an email about how to access the nodes. Use the information in the email in the next section to access your two end host nodes in WAIL.

Access to host nodes in your experiment is done via the Secure Shell tool. SSH is a standard secure remote terminal program and is available on all of the systems in the CSL labs. Other options are SecureCRT and PuTTY. In order to open GUI applications

you will need to be forwarding X11 packets (on by default in CSL). Windows machines will need an X windows system running. Try Cygwin. If possible work from the CSL labs, it will be much easier.

The login process to access systems in WAIL requires two steps. First, you must connect to a WAIL gateway system called ops.schooner.wail.wisc.edu. Next, from this computer you can connect to any of the experiment computers. To connect to the gateway computer type the command: `ssh username@ops.schooner.wail.wisc.edu` Substituting in your *username*. For example, type: `ssh bbadger@ops.schooner.wail.wisc.edu`

Answer yes to any authentication prompts. Enter your Schooner account password when prompted.

Read though the "New Experiment Started" email to find the Qualified Name of one of the two nodes in your experiment. This will be the destination of your next SSH connection from OPS.

To connect to one of the hosts type the command: `ssh -X QualifiedName` where Qualified Name is the name of computer you want to access. X11 forwarding is not enabled by default on OPS, so the -X flag tells SSH to forward X11 packets. (You can leave off the schooner.wail.wisc.edu part of the Qualified Name if you like)

For example: `ssh -X node1.test.cs640.schooner.wail.wisc.edu` Be sure to answer yes to any authentication prompts. Note: Connecting to the Qualified Name only works when connecting from inside the WAIL world. If you are attempting to connect from the outside world you will need to connect to OPS first.

You are now connected to one of the host computers in your simple network and can issue commands to the linux operating system running on the computer. You are encouraged to look around the system - see what applications and utilities are available. You are a superuser on this system so you can do whatever you like - even trash the OS configuration! Once you are finished with the node type exit to return to OPS. From there you can connect to other objects in your experiment or exit from OPS.

For security reasons you are only allowed access to the nodes in your experiment. Furthermore, commands that require superuser privileges will need to

be prefaced by the `sudo` command (type `man sudo` from the UNIX prompt to get help with `sudo`). This command executes the target command with root permissions. If you see any permissions errors try the command again with `sudo`. For example if you wanted to run the `tcpdump` utility, you would need to type `sudo tcpdump`.

You are also encouraged to look at the networking configuration on each system in your experiment (this is a networking lab after all!) using the `ifconfig` command. Each node will have (at least) two interfaces. `eth0` is used by Schooner to control the node. `eth1` is the emulated link to the LAN. In general you should never modify `eth0`. Changes to the configuration of `eth0` might cause you to lose your SSH connection to the node. But, go ahead and give it a try if you like.

PING is a standard utility to test connectivity between nodes. You should look at the manual page (via `man ping`) for PING if you are unfamiliar with it. Next, try to ping from one node to the other in your configuration. Do this by using the qualified name of the target for your probes. If ping is working correctly, you should see a line/response printed on the screen of the sending node approximately once per second. If it is not working, nothing will print and you will need to try to debug your use of the command. You can also try to make changes to the delays that you specified in your test configuration and then see what impact that has on the ping output (note, if you add e.g., a 10ms delay, your ping times should be on the order of 20ms - see if you can determine why this might be the case). You should also feel free to expand your configuration to include more systems or multi-hop paths - what happens when you try to ping in this environment? Finally, see if any standard applications (email or the web or others) are available on the hosts. If so, should you be able to run these applications in your experimental network? What happens when you run a utility like `tcpdump` when you are running a program like ping that generates traffic? Try it (note, you can open multiple terminal windows on a single system)!

When you are finished: Remember to close your experiment once you have finished with it. To close an experiment: Open the Schooner web page and log in. Open your "My Emulab" page by clicking on

```

----- test.report -----
Experiment: cs640/test
State: active

Virtual Node Info:
ID          Type      OS          Qualified Name
-----
node0       pc        RHL90-STD  node0.test.cs640.schooner.wail.wisc.edu
node1       pc        RHL90-STD  node1.test.cs640.schooner.wail.wisc.edu

Physical Node Mapping:
ID          Type      OS          Physical
-----
node0       i2        RHL90-STD  pc-i2-42-c
node1       i2        RHL90-STD  pc-i2-42-b

```

Figure 2: Qualified name listing sent via email, which is used to identify hosts in your experiment.

the link in the lower left corner of the screen. Click on the experiment name in "My Emulab", then select Terminate Experiment on the left hand side of the experiment frame. This will then release the resources you were using and make them immediately available to others.

3 Part 2: Building an "at home" LAN

This part of lab 1 will take you through some of the practical tasks associated with setting up a local area network (LAN) at home. It will include installing Ethernet cards in a host PC, configuration of Windows XP and a Linux end host, creation of a Cat 5 Ethernet cable and the use of a standard Ethernet switch.

3.1 Example Scenario

You have two personal computers at home and you would like to enable these two PC's to communicate and share an Internet uplink connection. You have two Ethernet network interface cards (NICs), a spool of category 5 cable, 4 RJ-45 connectors, a wire crimper, a simple switch, and an established Internet connection.

3.2 Sign Up and Lab Time

There will be a sign up sheet posted on the TA's office door to reserve times in WAIT for this part of lab 1. Please sign up for hour long slots so that teams do not bump into each other during this portion of lab 1. You shouldn't need more than two hours to complete this portion.

Please be sure to leave a few minutes at the end of your hour slots to return the equipment to the completely disassembled state in which you found it. This is important so that everyone can start from the same point. Please be courteous to your classmates - thanks!

3.3 Step 1: Cabling

There are two main types of physical cabling that are in wide use in networks today: optical fiber and copper wiring. There will be a number of different samples displayed during the lab, make sure to identify and record the types of wires in your notebooks.

3.3.1 Optical Fiber

Optical fiber has a core filament made from glass which transmits light, an example single-mode optical fiber is given in Figure 3.

In general, optical fiber has less attenuation (decrease in signal intensity) over large distances, and




Figure 3: The image shows the layering in a single mode optical fiber

can support higher bandwidths than copper. However, optical fiber has a complex manufacturing process, and can be costly and complicated to use due to the polishing and testing necessary when deploying cables.

3.3.2 Copper Wiring

Bundles of Category 5 copper wiring are widely used in fast Ethernet networking environments. Other examples of copper wiring in the networking sense include Cat 1-7 where the category specifies performance characteristics for that particular class of cable. If you are interested in the specification for Cat 5 cable see <http://www.dupont.com/cablingsolutions/pdf/h96496.pdf>

A Category 5 cable contains 4 twisted pairs of copper wire as shown in Figure 5.

How to make your own Ethernet cable from cat5 and RJ-45 connectors from http://www.ertyu.org/steven_nikkel/ethernetcables.html. You will need to make one cable for this part of the lab. You can use a lab-supplied Ethernet cable to connect your second PC to the switch.

1. Strip off about 2 inches of the cable sheath.
2. Untwist the pairs - don't untwist them beyond what you have exposed, the more untwisted ca-




Figure 4: The image shows the Category 5 cabling that you will be using to connect your PC's to the switch in part 2 of lab 1.

ble you have the worse the problems you can run into.

3. Align the colored wires according to the wiring diagram.
4. Trim all the wires to the same length, about 1/2" to 3/4" left exposed from the sheath.
5. Insert the wires into the RJ45 end - make sure each wire is fully inserted to the front of the RJ45 end and in the correct order. The sheath of the cable should extend into the RJ45 end by about 1/2" and will be held in place by the crimp.
6. Crimp the RJ45 end with the crimper tool.
7. Verify the wires ended up the right order and that the wires extend to the front of the RJ45 end and make good contact with the metal contacts in the RJ45 end.
8. Cut the cable to length - make sure it is more than long enough for your needs. Remember, an end to end connection should not extend more than 100m (328ft). Try to keep cables short, the longer the cable becomes the more it may affect performance, usually noticeable as a gradual decrease in speed and increase in latency.
9. Repeat the above steps for the second RJ45 end.

RJ45 - TAB Down




Figure 5: Wiring diagram for Category 5 twisted pair cable




Figure 6: PCI slots on a motherboard.

10. If a cable tester is available, use it to verify the proper connectivity of the cable (see Wiki on this).

3.4 Step 2: Network Interface Cards

There are many different types of network interface cards available. The standard differentiators are the backplane/motherboard interface type and the networking interface types. For this lab we are interested in PCI motherboard interfaces and Ethernet network interfaces since these are the two most commonly used today (do you know what other types of interfaces are still in use today?). You will be installing NICs into the two host PC's during this lab. You will be expected to follow the manufacturers recommendations for installing this hardware device on the motherboard. The procedure is provided by the Network Card Installation Reference given in the Reading section of the Pre-Lab. A picture of the PCI slots on a motherboard are given in Figure 6.

The procedure for installing a PCI/Ethernet NIC is as follows:

1. Shut down and unplug the computer.
2. Open the case.

3. Identify an empty PCI slot, and remove the back plate if needed.
4. Ground yourself to discharge static electricity which is harmful to the network interface card and motherboard by touching a metal portion of the computer chassis.
5. Carefully Seat the network interface card in the expansion slot, apply firm pressure on both ends of the card. When properly inserted the gold metal contacts of the card should barely be visible.
6. Anchor the card with the screw taken from the back plate.
7. Close the case, plug the computer into the power strip.
8. Restart the computer.

3.5 Step 3: Connecting via Switches

An Ethernet switch provides connectivity between the connected hosts. In addition, the switch allows multiple PC's to access the shared Internet uplink that we have provided for this lab. We will go into greater detail on switch operation in a later lab. The key point here is that small switches such as the one provided for this lab are plug-and-play devices that typically require no or very little configuration out of the box.

Plug one end of your Cat 5 Ethernet cables into the PCs and the other ends into the ports on the switch. Note, there is a special uplink port which should have a cable already installed. This enables access to the Internet. If things are working correctly, the small status lights on the switch should light up when you plug in the cables. If not, be sure the PC's are turned on. If this does not cause the status lights to illuminate, some debugging will need to be done - see the Wiki for ideas.

3.6 Step 4: Linux PC Networking Configuration

Power on the linux machine after the PCI card has been added. After the machine reaches a login prompt, login as user 'root' with password 'cs638'. Use the ifconfig utility to configure the IP address on the eth0 interface. The IP address information is listed on a sheet on the table with the PC's. Remember to bring up the interface with 'ifconfig eth0 up'. If your cabling is good and the interface properly configured, you should be able to ping the address 10.2.0.1, which is the default gateway of the network.

3.7 Step 5: Windows PC Networking Configuration

The Microsoft website (also provided under Readings in the Pre-Lab) has a nice set of reference materials and insights on networking configuration. See <http://support.microsoft.com/default.aspx?kbid=813936> for an 8 step tutorial on how to set up a home LAN with Windows XP. By following this tutorial, you should be able to communicate between PCs' on your LAN and with the greater Internet.

Try the ping utility between the two systems on your LAN by bringing up a terminal window on either system and using the IP address of the alternate system. Try accessing your favorite web site in the Internet from the Windows PC.

NOTE: Please turn in the cable that you made with your lab notebooks (obviously, only one lab partner can do this so please indicate who your partner was in your lab notebook if you do not turn in a cable).

3.8 Summary of LAN Setup Tasks

1. Observe and identify the equipment in the lab, ask the TA questions if you are unsure about any equipment or procedures.
2. Install the NIC PCI cards into the PC's.
3. Verify that Windows recognizes the NIC's by looking at the networking configuration window (you should see an Ethernet NIC listed).

4. Crimp together a single Cat 5 cable as directed above, the other cables will be provided.
5. Connect the PC's to the switch and see if a link lights are active.
6. Run ping between systems to see if they can communicate.
7. Access your favorite web site in the Internet from both PC's.