



White Paper

Methodology for Testing Wireless LAN Performance

Methodology for Testing Wireless LAN Performance

Introduction

Whether you want to evaluate the performance of wireless LANs (WLANs) in an informal way or through precise benchmarking procedures, the first step is to understand the factors involved. The ease of setting up and using WLANs makes it easy to overlook many crucial factors and their resulting performance variations. These performance variations can be extreme, however, and they make all the difference in the cost, security and viability of a wireless network.

This white paper describes the factors that affect a WLAN's throughput and coverage, then provides a detailed methodology for those who want to benchmark throughput and coverage in a disciplined way.

An Overview of Throughput and Coverage Factors

A WLAN generally consists of an access point (AP) that connects to a wired network and remote devices (client) that connect to the access point through wireless (radio) links. Throughput is defined as the speed with which a user can send and receive data between a remote device and the access point. Throughput varies across the WLAN's coverage area. This section profiles the main factors that determine WLAN throughput and coverage.

1. 802.11 Protocol—The IEEE 802.11 standard defines various physical-layer rates for different types of WLANs, such as 1, 2, 5.5 and 11 Mbps for 802.11b and 802.11g. Rates for 802.11a and 802.11g include 6, 9, 12, 18, 24, 36, 48 and 54 Mbps. The user throughput is less than these link rates for several reasons:

- Each packet includes additional data, such as preambles, headers (MAC, IP, TCP, etc.) and checksums.
- When every directed (unicast) packet is received, the receiver transmits a short acknowledge packet back to the sender.

- Transmitters wait for short random times between packets to allow other users to contend for and share the channel.

Given these reasons, the theoretical maximum user-level performance for the various 802.11 systems is:

Table 1-1. Maximum Theoretical Application-Level Throughput

	Number of Channels	Modulation	Maximum Link Rate	Maximum TCP Rate	Maximum UDP Rate
802.11b	3	CCK	11 Mbps	5.9 Mbps	7.1 Mbps
802.11g (with 11b)	3	OFDM/CCK	54 Mbps	14.4 Mbps	19.5 Mbps
802.11g (11g-only mode)	3	OFDM/CCK	54 Mbps	24.4 Mbps	30.5 Mbps
802.11a	19 ¹	OFDM	54 Mbps	24.4 Mbps	30.5 Mbps
802.11a TURBO™	6	OFDM	108 Mbps	42.9 Mbps	54.8 Mbps

- ¹ 13 non-overlapping channels in the United States and up to 19 non-overlapping channels in Europe depending on local regulations.

[Table 1-1](#) assumes 1500-byte packets, encryption enabled, default 802.11 MAC configurations, zero packet errors, and maximum available channel bandwidth (that is, operating at close range). Note that some 802.11 implementations use tricks such as reducing backoff times between packets to improve throughput performance. Such tricks can result in interoperability problems with other vendors' systems.

[Table 1-1](#) also shows two rates for 802.11g to account for the lower rates in 802.11b compatibility mode. The throughput of an 802.11g WLAN decreases significantly in 802.11b compatibility mode because every 802.11g (OFDM) packet needs to be preceded by a CTS packet exchange recognizable by legacy 802.11b devices. With no 802.11b devices connected, an 802.11g network can operate in 11g-only mode and should achieve the standard throughput of 802.11a. The current 802.11g draft standard also provides for a slower RTS/CTS header (instead of CTS-only) when in 802.11b compatibility mode, which will further reduce the 14.4 Mbps TCP/IP rate to 11.8 Mbps.

You therefore have two choices with 802.11g networks: You can achieve high rates comparable with those of 802.11a networks. Or you can get 802.11b compatibility. You can not have both at the same time.

Since the key feature of 802.11g is backward compatibility with 802.11b, throughput tests should be done with an 802.11b client device connected to the access point but otherwise idle. This setup ensures that the 802.11g network is operating in an 802.11b compatible mode.

2. The radio environment—Several issues affect the way the radio signal travels from one device to another:

- Radio energy attenuates when it propagates. As radio waves propagate outwards spherically, the energy spreads over an ever-increasing area. In free space, doubling the distance decreases the received power by a factor of 4—the so-called $1/r^2$ behavior.

Radio signals also attenuate when they pass near or through objects such as floors, walls, furniture and people. The attenuation increases with the object's conductivity (due to metal or water content, for example). The combination of these two attenuation effects reduces radio signal strength by $1/r^3$ to $1/r^4$, or even $1/r^5$. In other words, each time you double the distance, the received power might decrease by 8 to 16 times.

- Antenna designs affect how much radio-frequency (RF) energy is transmitted or received and where it is directed.
- Scattering and multipath cause fading effects. Signal strength can change rapidly as a function of location because the received signal is the sum of potentially numerous signals scattered from nearby objects. As the transmitter or other objects in the environment move, the scattered signals sometimes add together and sometimes cancel each other. Fading can change significantly over distances of a wavelength or so (12.5 cm at 2.4 GHz and 6 cm at 5 GHz).

Fading also occurs over time as well as location. Even small changes in the environment (for example, people or other objects moving) can affect the fading pattern. This means that the received signal strength can also change quite quickly over time, even when the receiver and transmitter are fixed.

- Scattering and multipath results in delay spread. The received signal might contain several slightly delayed copies of the transmitted signal, as the scattered signals travel via different physical paths of different lengths.
- Other devices occupying the same or nearby channels cause interference. For example, the 2.4 GHz spectrum might be occupied by Bluetooth devices, microwave ovens, and cordless telephones.

3. Frequency—A common misconception is that free-space propagation depends upon frequency, so higher frequencies are assumed to propagate less well than lower frequencies. As a good counter example to this misconception, consider visible light, which is simply an ultra-high frequency electromagnetic wave that propagates perfectly well across large distances.

On the other hand, effects such as antenna efficiency, RF component performance, and absorption through and scattering around objects do depend upon frequency. Here are some of the frequency-dependent effects:

- Generally, antennae of the same physical size tend to become more directional (have higher gain in some directions and less in others) as the frequency increases.
Advantage: 5 GHz.
- Absorption due to propagation through objects tends to increase with frequency.
Advantage: 2.4 GHz.
- Scattering around objects might have a positive or negative effect on signal strength as a function of frequency, depending upon the relative sizes and locations of the objects.
Advantage: Neutral.

- Noise and spurs generated by nearby electronics (for example, inside the AP or PC laptop) in addition to co-channel interference, such as Bluetooth devices, cordless phones and microwave ovens, will degrade 2.4 GHz sensitivity more than 5 GHz.
Advantage: 5 GHz.
- Cable loss increases with frequency, so antenna cables (if present) in the AP or laptop will have more loss at high frequency, unless more expensive cables are used.
Advantage: 2.4 GHz.

In more open environments, there will be little difference between 2.4 GHz and 5 GHz propagation. For example, measurements of 2.4 GHz and 5 GHz propagation done by WJ Communications in two indoor environments show little difference between 2.4 GHz and 5 GHz propagation. See the full paper at http://www.watkins-johnson.com/pdf/techpubs/Indoor_prop_and_80211.pdf

Typically, the OFDM modes of 2.4 GHz 802.11g networks will have slightly less coverage than 2.4 GHz 802.11b networks. Depending upon the propagation environment, the coverage of 5 GHz 802.11a networks might be similar to, or in some cases less than, that of 802.11g networks. The differences between 2.4 and 5 GHz propagation are generally insignificant compared to the differences between one vendor's equipment and another's, however. An 802.11a product from one vendor might have better coverage than an 802.11g product from another vendor.

4. The vendor equipment design—Equipment from different vendors exhibit significantly different performance due to architecture, design, manufacturing and software variations, as well as proprietary features and enhancements.

5. Vendor interoperability—Products that undergo Wi-Fi certification are certified to interoperate with a wide variety of vendors' products. However, these tests mainly verify basic connectivity and do not enforce stringent throughput requirements. You might be able to connect a client device to a different vendor's access point, but you might not be getting very high throughput. Products that provide good performance (throughput, coverage, etc.) when connected to a variety of different vendor's devices are clearly more desirable.

6. Security—Security includes encryption and authentication. Encryption protects WLAN traffic from eavesdropping and other attacks such as replay or man-in-the-middle attacks. Authentication validates the users' credentials (ensuring that the user is who they say they are) and also possibly validates the network's credentials (ensuring that the network is what it says it is, and not someone masquerading as the network).

WLAN security standards have progressed from WEP to TKIP and WPA and now to AES (the Advanced Encryption Standard), with significant security enhancements at each stage. No matter what security standard is involved, the way the standard is implemented can affect the WLAN's performance. Specifically, some vendors implement encryption in software, which can dramatically reduce throughput compared to advertised rates. When evaluating performance, it is vital to measure throughput with encryption enabled. For more details, see http://www.atheros.com/pt/atheros_wlansecurity.pdf.

Measuring Throughput and Coverage

The throughput of WLANs depends heavily on the environment, including the distance between the client and the access point. The throughput generally falls off as distance increases, but factors such as obstructions (like furniture, people, or walls of different construction) also have a significant effect. Throughput does not depend upon distance alone. It is possible to have distant test locations that produce higher data rates than closer locations. Moreover, the peak data rate measured at short distances is not the most important factor in the user's experience. Rather, the rate the user experiences at a variety of distances and locations is a very important factor. Therefore, it is critical to measure WLAN throughput at a variety of locations, including some far from the access point.

WLAN environments generally fall into three categories:

- **Outdoor:** typically a direct line of sight between the access point and client. Examples include outdoor campus coverage, public areas, or even inside large, open buildings such as airport concourses or convention halls.
- **Open office:** no longer a direct line of sight between the access point and client, but typically at most two-to-three obstructions such as walls. Examples are warehouses or offices containing cubicles, lobbies and meeting areas.
- **Closed office:** no direct line of sight, with many obstructions between the access point and the client. Examples are buildings with regular offices and many walls.

WLAN coverage differs significantly in these different environments. Outdoor WLANs provide the longest ranges and closed-office WLANs the shortest. Different construction techniques also have a significant impact on coverage and throughput. For instance, concrete walls attenuate signals more than stud walls with sheet rock.

In general, the relative performance and throughput for different products under test should be similar across the different environments. So if Vendor #1's product is significantly better than Vendor #2's in an open-office environment, it is highly likely (although not guaranteed) that it will be significantly better in other environments. It is possible (although more time consuming) to test products across several different environments to accurately determine the relative performance.

Various benchmarking programs can be used to measure the throughput the user will experience. Typically these measure TCP throughput in megabits per second (Mbps) in either the uplink direction (for example, upload from the client to the AP) or downlink direction (for example, download from the AP to the client). Downlink TCP performance is the most relevant metric, since it reflects the most common usage such as browsing the web or downloading email.

Some applications like video streaming use a simpler protocol called UDP. Generally, UDP performance numbers will be 15-20 percent higher than TCP performance numbers because there is less protocol overhead associated with UDP. Many of the benchmarking programs can measure UDP performance, too.

One of the following benchmarking programs should be used:

- Chariot from NetIQ: a powerful commercial tool for network testing. Chariot allows complex traffic patterns to be modeled. See <http://www.netiq.com/support/chr/specs.asp>.
- Qcheck from NetIQ: a simple (and free) utility for network benchmarking for Windows platforms. See <http://www.netiq.com/qcheck/default.asp>.
- FTP: a widely available protocol for downloading and uploading files. One drawback is that FTP also has to read/write data from the disks, which can affect measurements that would otherwise have high throughput.
- Netperf: a free, open-source benchmarking program originally from HP Labs. Netperf supports all flavors of Unix/Linux and Windows platforms. See <http://www.netperf.org/>.

Refer to Appendix A for details on test setup and Appendix B for details on the correlation of the results gathered from using the various benchmarking programs.

Test Setup

The first step is to decide which products will be tested and which access points will be tested with which client cards. The natural test configuration is to pair the access point and client from the same vendor. However, it is also important to assess how well a client card performs with other vendor's access points, since many users will use their client devices in many different networks.

An example of test configurations with natural pairing:

- Test 1: Vendor 1 access point with Vendor 1 client.
- Test 2: Vendor 2 access point with Vendor 2 client.
- Test 3: Vendor 3 access point with Vendor 3 client.

An example of test configurations where just client devices are being tested against a 3rd-party access point:

- Test 1: Vendor 4 access point with Vendor 1 client.
- Test 2: Vendor 4 access point with Vendor 2 client.
- Test 3: Vendor 4 access point with Vendor 3 client.

An example of test configurations where interoperability is being tested:

- Test 1: Vendor 1 access point with Vendor 1 client.
- Test 2: Vendor 1 access point with Vendor 2 client.

- Test 3: Vendor 2 access point with Vendor 1 client.
- Test 4: Vendor 2 access point with Vendor 2 client.

The second step in designing the test procedure is to choose a set of test locations:

- Select a test location for the access point. Ideally the access point should be located high above the floor and away from immediate obstructions. Most importantly, use exactly the same access point location for each product tested.
- Select a channel for testing, and verify that the RF environment on the selected channel is clear. Use a sniffer or client device to check that there are no access points or ad-hoc networks located on the same channel throughout the test area. For 11b and 11g, this means no overlapping channel; channels with number spacing of 4 or less overlap and cause significant in-band interference. For example, 2.4 GHz channel 1 overlaps with channels 2, 3, 4, 5, and channel 6 overlaps with channels 2, 3, 4, 5, 7, 8, 9 and 10. For 11a the standard 54 Mbps channels do not overlap.
- Select at least eight test locations at a variety of locations and distances from the access point (see [Figure 1-1](#)). At least one test location should be at the limit of coverage. (If you later discover that one product under test has much better coverage than initially expected, then additional, more remote, test locations need to be added and the earlier tests with the other equipment to be repeated at these new locations.)
- All wireless LANs have a limit on signals that are too strong. Some WLAN products may actually produce low data rates at very close ranges (for example, less than 3 feet). Therefore, the closest test points should be no less than 5 feet apart.

The key criterion is repeatability. For each product under test, the access point locations, software setup, channel used, overall environment, test procedure and test locations should be the same. Environmental repeatability is generally improved if the tests are done back-to-back, for example, over as short an elapsed time as possible.




Figure 1-1. Typical Range and Throughput Setup

At each location, make a minimum of three measurements so you can average-out some of the local radio fading effects by repeating a measurement with a small shift in the test laptop's location. This strategy improves the test's repeatability and makes the results less prone to "lucky" or "unlucky" measurements each time the test is repeated.

Thus, at each test location, repeat at least three times:

- Measure the downlink and uplink TCP (and optionally UDP) throughput.
- Displace the test laptop by about one wavelength (that is, between 12.5 cm at 2.4 GHz and 6 cm at 5 GHz) and repeat. Alternatively (or additionally), rotate the laptop by 45 degrees or more.

As different products are tested, use the identical procedure (test setup, software, locations and displacements or rotations).

Test Procedure

Putting all the previous steps together, the overall test procedure is:

1. Setup test #1: Install the access point from Vendor #1, and client card #1.
2. Go to the first test location and make at least three measurements, moving the location and/or orientation of the laptop slightly between measurements.

3. For each location, record the throughput (TCP downlink, TCP uplink, and optionally UDP downlink and uplink) for each run. Record the test locations on a floorplan.
4. Repeat steps 2-3 for each test location.
5. Repeat steps 1-4 for each different equipment configuration.

Performance Metric

A combined metric, or score, of throughput and coverage can be computed from the measurements. Both high throughputs and large coverage areas are desirable.

The metric should be proportional to the measured throughput. For example, if one product produced exactly twice the throughput at each location compared to a second product, its metric is twice as large.

The metric should also be proportional to the coverage area. For example, if one product provides, say, 10 Mbps over a certain area (and no connection outside this area), while a second product provides 10 Mbps over twice the area (and no connection outside), then the score of the second product is twice the first.

To compute this metric, compute the average throughput of the three (or more) measurements at each test location. Also, compute or measure the straight-line distance in meters from each test location to the access point.

The performance metric is the sum over locations of the throughput Mbps(i) multiplied by the ring area over which that throughput is achieved:

$$\text{Performance Metric} = 0.001 \times \sum_{i=1}^n \text{Mbps}(i) \left[r^2(i) - r^2(i-1) \right]$$

n is the number of locations at which measurements were made. The range r(i) is squared because area is proportional to the square of the distance. The normalization factor 0.001 is included to make the end figures more tractable. The area of the previous range r(i-1) is subtracted from the current range to obtain the ring area. As a rule, r(0) is the origin and has zero value.




Figure 1-2. Same Throughput Over Different Coverage Radii

As a first example, consider two access points, AP1 and AP2 (see Figure 1-2). Both provide 20 Mbps coverage up to a certain point and drop off immediately after that point. In this case, there is only one measurement location, so $n = 1$. Assume that the range of coverage for AP1 is 10 meters, while that for AP2 is 14.1 meters. The performance metric for AP1 is 2.0, while that for AP2 is 4.0. This metric is consistent with the fact that AP2 provides twice the coverage area as does the AP1.




Figure 1-3. Different Throughput Over Different Coverage Radii

Consider a second example in which AP1 provides higher average throughput at close range but smaller overall coverage area than AP2 does (see Figure 1-3). Specifically, the average throughput at the various ranges are as follows:

Table 1-2. Data for Example 2

Measurement Location (m)	AP1 Throughput (Mbps)	AP2 Throughput (Mbps)
$r(1) = 30$	20	18

Table 1-2. Data for Example 2

Measurement Location (m)	AP1 Throughput (Mbps)	AP2 Throughput (Mbps)
r(2) = 70	15	13
r(3) = 100	0	10

The performance metric for AP1 is $0.001 ([20(30^2 - 0) + 15(70^2 - 30^2) + 0(100^2 - 70^2)]) = 78$. The performance metric for AP2 is $0.001 ([18(30^2 - 0) + 13(70^2 - 30^2) + 10(100^2 - 70^2)]) = 119.2$.

Thus, the performance metric is truly a function of both average throughput and range. At the far edges of coverage, extra range contributes greatly to the performance metric, due to the squared dependency.

What the performance metric is truly quantifying is system capacity, that is, a system's ability to deliver high data rates across as wide an area as possible. In the above example, AP2 delivers far greater system capacity because of its extended range of coverage, even though its average data rates are lower than those of AP1.

Appendix A: Test Setup

A common step that should precede any of the steps outlined below is a wired throughput check. A crossover Ethernet cable should be used to connect the server PC and the laptop being used for range testing. In this way, the performance of the peripheral busses on the PCs, as well as the functionality of the test programs, can be verified. A 100-Mbps Ethernet connection should register TCP throughput of 85-90 Mbps in this configuration. Verification of this increases the confidence that the performance of the wireless link is not affected by host hardware issues.

Furthermore, the server PC should reside on an independent subnet or network from the corporate network. This ensures that the server PC-to-AP connection is not affected by traffic outside of the test setup.

FTP setup is not discussed in this document because implementations of FTP clients and daemons vary too widely to make generalization possible. While FTP throughput performance is useful as a rough indicator of network performance, it does not provide the degree of accuracy and repeatability that dedicated measurement tools do.

Chariot and Chariot Lite Setup

NetIQ Chariot enables the creation of complex scripts that test various network performance dimensions (throughput, transaction rate, response time) across a range of protocols (TCP, UDP, RTP, IPX). NetIQ Chariot Lite application enables the execution of Chariot test scripts but not the creation nor modification of them. To do so requires a full Chariot license.

Setup of the Chariot test environment requires the installation of a NetIQ Endpoint at the server end and Chariot or Chariot Lite at the client end. NetIQ Endpoint can be downloaded from <http://www.netiq.com/support/pe/pe.asp>. Once the Endpoint has been set up, only the IP addresses of the Endpoint and client need to be entered into the Chariot environment for test runs to begin.

One of the key features of the Chariot environment is the ability to view the time series of the measurement parameter under consideration. An example screen shot capturing throughput performance is shown in [Figure 1-4](#).

Chariot and Chariot Lite ship with default test scripts denoted by the .scr suffix. They serve as the starting blocks with which to build customized test scripts. In particular, the scripts `Filercvl.scr` and `Filesndl.scr` are useful for testing sustained throughput performance with large amounts of data transfer.

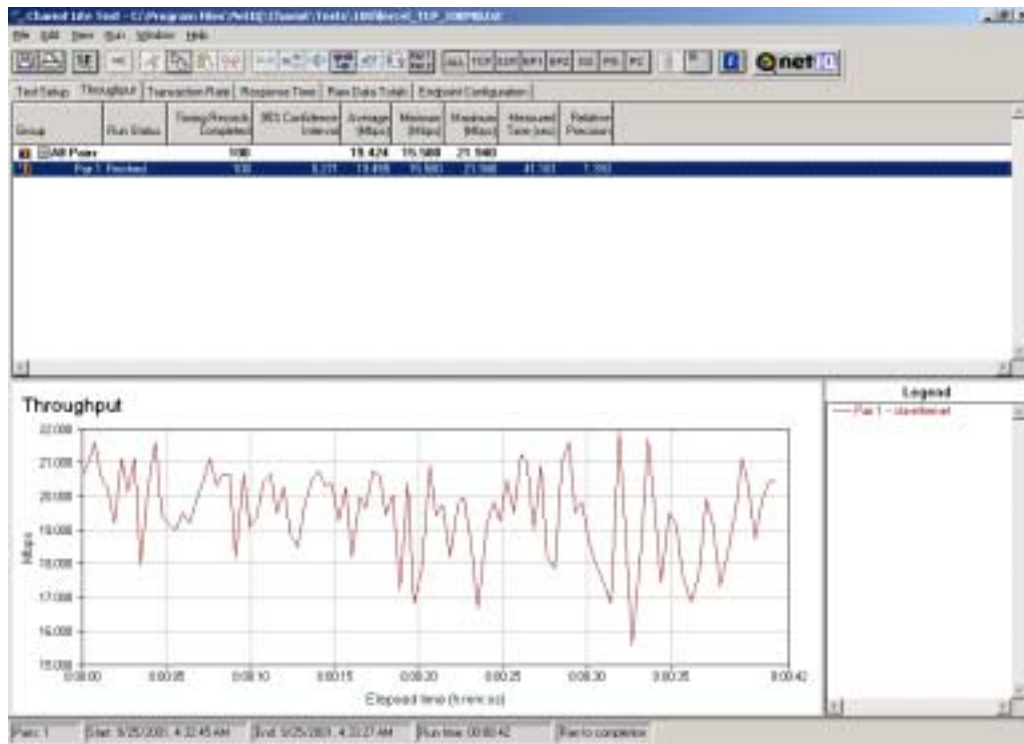


Figure 1-4. NetIQ Chariot Lite User Interface

QCheck Setup

The QCheck application is freely downloadable from NetIQ's website at <http://www.netiq.com/qcheck/default.asp>. Once installed, NetIQ presents the user with a simple interface, shown in Figure 1-5.

QCheck requires a NetIQ Endpoint to be installed on the server end, much like Chariot. NetIQ Endpoint can be downloaded from <http://www.netiq.com/support/pe/pe.asp>, or alternatively QCheck can be run on the server machine. The most useful tests for the purposes of range and throughput measurement are "Throughput" and "Streaming." However, the streaming data rate is limited to 1 Mbps and the test duration is bounded at 30 seconds, which limits the applicability of QCheck to simple test cases.

The test protocol should be set to "TCP" (or "UDP" if desired). For an uplink test the "From Endpoint 1" should be set to "localhost" and the "To Endpoint 2" should be set to the IP address of the server machine connected to the access point (on which the Endpoint is running). For a downlink test, the test needs to be run from QCheck running on the server machine with "From Endpoint 1" set to "localhost" and "To Endpoint 2" set to the IP address of the client (remote) machine. The Data Size should be set to 1000 kBytes (the maximum value). The most useful test for the purposes of range and throughput measurement is "Throughput". Each time "Run" is selected, this transfers 1000 kBytes using TCP and reports the throughput. The "Streaming" test is limited to 1 Mbps and the test duration is bounded at 30 seconds, which limits the applicability of QCheck to simple test cases.

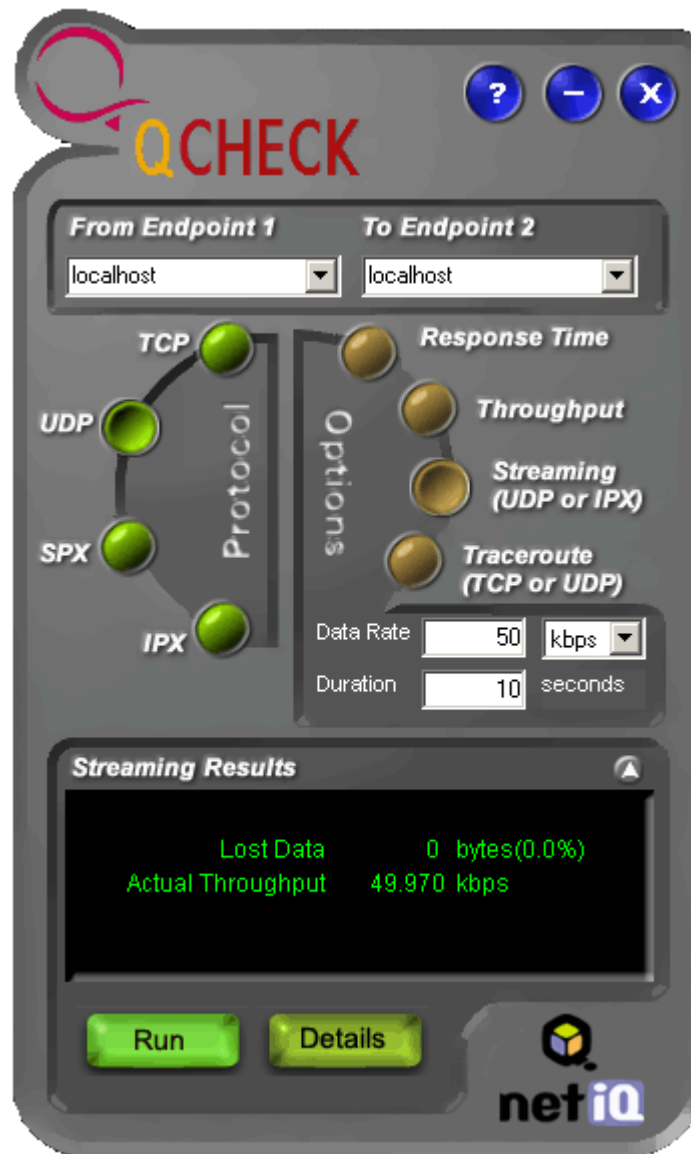


Figure 1-5. Net IQ QCheck User Interface

Netperf Setup

To enable bi-directional testing, install netperf.exe and netserver.exe on both the server PC and the client station. *Note: Netserver must be started on the host PC that is being targeted at by the PC running Netperf.*

The typical syntax for a netperf command used for range and throughput testing is:

```
netperf -N -H 192.168.1.1 -l 10 -t UDP_STREAM -- -m 1470
```

This specifies an uplink (client station to server PC) UDP test of duration 10 seconds, with send packet size 1470 bytes. The 1470 byte packet size is chosen so as to not cause UDP to fragment the packet. 192.168.1.1 is the server PC's IP address. For a TCP test, specify TCP_STREAM for the -t flag. The full command is as follows:

```
netperf -N -H 192.168.1.1 -l 10 -t TCP_STREAM
```

To perform downlink (server PC to client station) testing remotely from the client station, run an rsh daemon on the server PC. This enables remote command line interaction with the server PC. Note: installing the rsh daemon with no password allows any remote station with knowledge of the server PC's IP address to have unimpeded browse, edit and execution access, posing a serious security risk. The rsh daemon should only be run on test setups isolated from the corporate network and it should be stopped once the testing is complete.

To perform a downlink test, the command is:

```
rsh 192.168.1.1 netperf -N -H 192.168.1.5 -l 10 -t UDP_STREAM -- -m 1470
```

This specifies a downlink (server PC to client station) UDP test of duration 10 seconds, with send packet size 1470 bytes. The server PC has IP address 192.168.1.1, while the client station has IP address 192.168.1.5. For a TCP test, specify TCP_STREAM for the -t flag. The full command is as follows:

```
rsh 192.168.1.1 netperf -N -H 192.168.1.5 -l 10 -t TCP_STREAM
```

Special consideration should be given to the send packet size. The Ethernet MTU is 1500 bytes, which means that any packet larger than that will be fragmented. The specified packet size should thus not exceed 1472 bytes so that when IP and TCP headers are included, the overall packet does not exceed 1500 bytes in length.

All of the binary executables necessary for Netperf testing are packaged in an archive file (Netperf_benchmarking_tool.zip) and the file is available for download at <http://www.atheros.com/pt/papers.html>.

Appendix B: Correlation of Results Across Platforms

A common question among testers is what degree of correlation exists across the throughput results given by the various test platforms. To address that concern, representative measurements are shown in [Table 1-3](#), [Table 1-4](#), and [Table 1-5](#) for three randomly selected locations at Atheros' Sunnyvale offices.

Table 1-3. Location 1

Application	TCP Downlink (Mbps)	TCP Uplink (Mbps)
Netperf	22.41	22.27
Chariot Lite	22.96	22.85
QCheck	22.24	22.22

Table 1-4. Location 2

Application	TCP Downlink (Mbps)	TCP Uplink (Mbps)
Netperf	16.75	17.54
Chariot Lite	18.79	16.71
QCheck	17.50	16.00

Table 1-5. Location 3

Application	TCP Downlink (Mbps)	TCP Uplink (Mbps)
Netperf	13.04	11.09
Chariot Lite	12.17	12.07
QCheck	14.68	11.27

As is evident from the data, there are no statistically significant discrepancies between the test platforms. The observed standard deviations between successive measurements on the same platform, on the order of 1.5-2.0 Mbps for test locations with weak signals, were greater than the differences between platforms.

The choice of test platform thus rests upon the richness of the feature set surrounding the core throughput measurement capability. There do not exist significant differences between popular platforms with regard to throughput measurement capability.

NOTE: This document may be updated periodically. Please check <http://www.atheros.com/pt/papers.html> for the latest versions of this and other Atheros white papers.

The information in this document has been carefully reviewed and is believed to be accurate. Nonetheless, this document is subject to change without notice. Atheros assumes no responsibility for any inaccuracies that may be contained in this document, and makes no commitment to update or to keep current the contained information, or to notify a person or organization of any updates. Atheros reserves the right to make changes, at any time, in order to improve reliability, function or design and to attempt to supply the best product possible.

Atheros and the Atheros logo are registered trademarks of Atheros Communication, Inc. All other trademarks mentioned in this document are the property of their respective owners.

Document Number: 991-00003-003



**Atheros Communications
Incorporated**
529 Almanor Avenue
Sunnyvale, CA 94086
t: 408/773-5200
f: 408/773-9940
www.atheros.com

Subject to Change without Notice