



# Internet Download Testing

*This application note describes Internet Download testing on the Trend Aurora Presto DSL Service Tester.*



*Application Note ANP INTDLD-18*

A large, light blue network diagram with several circular nodes connected by lines, serving as a background for the 'xDSL' text.

# xDSL



---

## APPLICATION LAYER TESTING

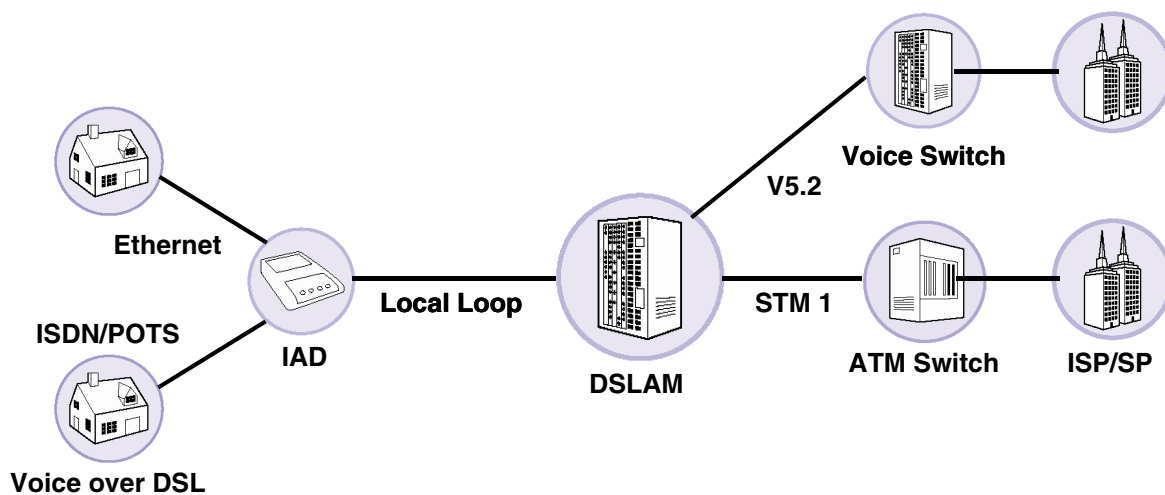
### Internet Download (HTTP)

This application note describes Internet Download testing using the Trend Aurora Presto DSL Service Tester.

The Internet Download option enables the user to connect to any valid web site using the HTTP (HyperText Transfer Protocol) Protocol and then download the HTML content of the page.

Data presentation is in the form of download statistics enabling the user to evaluate the performance at the application level.

For the purposes of this document, a typical DSL deployment is assumed, as shown in Figure 1.



---

Figure 1 Typical DSL Deployment.

### Limitations on Testing Download Performance

There are a significant number of factors that can contribute to the speed of a download from the Internet. Consideration of these factors is important when testing download speeds.

#### Network Routing

Like any major transport system, the Internet has many different routes through it, each with their own capacity and speed restrictions. Data sent to a computer has to wait while other data passes through routers, the Internet equivalent of a junction, on its way to the PC.

The connection to the ISP is just one part of a system that gets data to the PC or Tester. Between the tester and the server there are probably dozens of routers, communications links, and other network components that the data has to travel through. Each of these components has a fixed capacity and speed at which they can operate, and most of them handle network traffic for thousands, even millions of computers every day.

The inevitable result is that all Internet traffic encounters some sort of delay as it travels across the Internet, and this delay reduces the amount of data that is transmitted into the local DSL connection.

### **Traffic Density**

Slow Internet traffic is more likely to be encountered during peak use hours than during those times when fewer people are online, such as after 11 p.m. and before 7 a.m. Also, the traffic conditions at the destination may affect the download speed (for example, when testing American sites, better speeds will be obtained before lunchtime when the East Coast wakes up and goes online). Optimal results are obtained around 3:00am (GMT) on a weeknight, when almost everyone is asleep and Internet use is at its lowest.

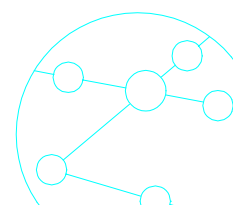
### **Redundancy**

No Internet connection ever performs at 100%. Controlled testing has shown a 33% loss of rated speed on the simplest possible network.

### **Multiplexing**

Whilst the customer DSL connection may be a dedicated line, all of an ISP's connections (DSL, ISDN, Dial Up, etc) are combined into one or more shared connections. In most cases, these shared connections have less capacity than the combined total of all the customer connections they serve - this is called the contention ratio, and ISPs can easily overbook their main Internet connectivity by ten, twenty or even fifty times.

When this is done with care it works extremely well. Since most Internet users spend more time reading their email and web pages than they do downloading them, they are only using a fraction of their connection's actual capacity. Overbooking allows an ISP to combine several customer connections into a single link that is smaller (and less expensive) than the combined total of all the connections they serve; without reducing the amount of data sent to a customer when they are downloading data.



The problem is that some ISPs, cable modem companies, and DSL providers can funnel so many connections into a small combined connection that normal customer demand overwhelms the capacity of the combined connection. This is a particular problem during peak use hours, when line speeds are significantly reduced.

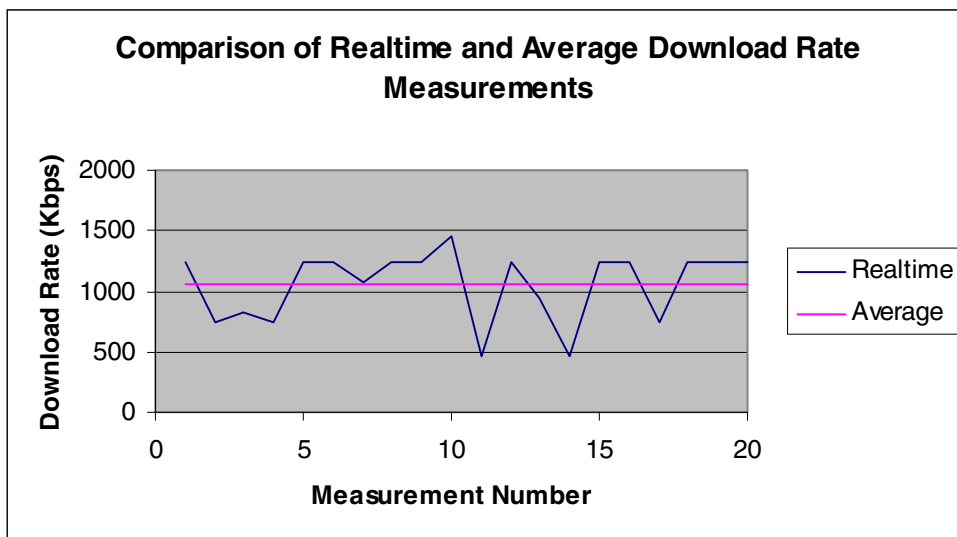
### Internet Download Test Methods

#### Single/Multiple Page Downloads

The HTML content of web pages is small - typically between 2 and 8k. When such a small amount of data (**a single web page**) is downloaded, the results can be highly variable depending on environmental factors such as those described above.

A simple way to overcome this problem is to download the web page repeatedly so that the 'peaks' and 'troughs' in download rate are averaged out. The measurement provided from this **'multiple download'** test is more representative of the download rates a user can expect to achieve during a typical Internet session.


The graph below (Figure 2), shows a plot of actual download rates (realtime) measured at 10 second intervals, compared with an average of the realtime data rate. The tests were conducted on a 2Mb broadband link and demonstrate the variability of individual downloads.



Figure

2

Comparison of Realtime and Average Download Rate Measurements.



Note from the graph that whilst a 2Mb link was available, the peak rate achieved during this test was around 1.4 Mb. (or 70% of the 'available bandwidth').

Good test practice would be to utilise the 'Single Page' and 'Multiple Page' download methods as follows:

**Single Page:** Use to validate connection to a web site or for fault finding using HTTP error codes generated through connection failure.

**Multiple Page:** Use to indicate average download rates from a particular connection/ web server. Can also be used for fault finding using HTTP error codes generated through connection failure.

## **Internet Download Results Interpretation**

From a user perspective, downloading a web page is a single, simple step that includes entering the URL (or selecting the required link) and then waiting for the page to download.

In reality, the download consists of a number of steps that can contribute to the total download time. These steps can broadly be defined as follows:

1. DNS Address resolution
2. Web server Response
3. Download from web server


The effect of these individual steps is discussed in the following sections.

### **Latency**

The domain name system (DNS) is the way that Internet URLs (www.....) are located and translated into IP addresses. A URL is a meaningful and easy-to-remember identity for an Internet address.

A DNS server is responsible for translating the URL into an IP address. The DNS server can be responsible for delay when downloading a web page for a number of reasons. Delay in download which can be attributed to the DNS server is known as 'latency' and is defined as the time between the start of the test and when the first HTTP response message is received from the DNS server.

In order to reduce the length of time of name resolution, and to reduce traffic on the network, an important concept of DNS is that of caching. Whenever a name server receives address information for another host or domain, it stores that information for a specified period of time (as determined by the Time-To-Live (TTL) value).



In this way, if another name resolution request for that host or domain is received, the name server has the address information ready, and does not need to send another address enquiry across the Internet.

The ability of the DNS server to cache resolved URLs should be considered when testing download rates.

Latency from the DNS server can be a significant contributor to low download rates and should always be measured independently from the 'Response' and 'Download' times/rates.

#### **Response Statistics**

Once the IP address of the web page has been resolved by the DNS server and then communicated back to the PC or tester, a request to download can be sent to the web server.

Depending on the volume of requests that the web server is currently handling, it can respond to the download request in anything from <1 second to >1 minute. This is the 'Response Time' and can contribute significantly to the overall test time.

#### **Download Statistics**

This is the total time taken to download the specified web page after the web server has responded to the download requests.

Also provided are the number of bits and bytes downloaded from the web site. This figure includes the HTTP header, transmitted at the start of every download.

