

HP Internet Reporter

Application Note 1274

LAN and WAN baselining and benchmarking

Summary Statistics Chart Builder

Selected Nodes/Stations Profiles (fps)

- #10 - 3D Area
- #11 - 3D Area w/Other
- #12 - Stacked Area
- #13 - Stacked Area w/Other
- #14 - Relative % w/Other
- #15 - 3D Pie

Selected Protocols Profiles (fps)

- #20 - 3D Area
- #21 - 3D Area w/Other
- #22 - Stacked Area
- #23 - Stacked Area w/Other
- #24 - Relative % w/Other

General Utilization

- #30 - Util(%)
- #31 - Util(kB)
- #32 - Broadc
- #33 - Util(%)
- #35 - FDDI F

Reset Selections

Microsoft Excel - ENDEMO.RWS

Date		General Stats				Errors (per sec)				Misc Stats			Selected Protocols (per sec)					Selected Nodes (per sec)											
Time	Utilization (%)	Utilization (fps)	Utilization (kbps)	Bytes Per Frame	Active Node Count	Total Errors	Jabbers	Bad FCS's	Runts	Collisions & Jams	Misaligns	Broadcasts (fps)	Broadcasts (%)	Multicasts (fps)	IP DOD	DECnet IV	XNS IDP	DEC LAT	Other Protocols	DEC--000CDC	Recall001E81	HP--002816	Other Nodes	DEC--000704	9-com--00427	UB--57D100	Recall002814		
17-May-91 15:36:49	6	114	73	641	5	3	1	1	1			8	7	10	24	53	24	8	5	26	24	24							
15:36:50	9	167	116	692	7							18	11	12	52	66	25	18	6	32	24	51	1	18	40	1			
15:36:51	11	186	135	729	5							22	12	10	66	71	22	22	5	32	22	66							
15:36:52	11	188	136	722	5							22	12	12	66	72	22	22	6	34	22	66							
15:36:53	11	326	142	436	7	3	1	1	1			58	18	10	60	112	91	58	5	83	30	59	1	58	34	1			
15:36:54	11	385	133	345	5							75	19	10	50	130	125	75	5	110	125	50							
15:36:55	11	387	133	344	5							75	19	12	50	131	125	75	5	112	125	50							
15:36:56	26	615	322	524	7	3	1	1	1			183	30	10	143	174	109	183	6	77	108	142	2	183	102	1			
15:36:57	35	759	436	574	5							250	33	12	197	206	100	250	6	62	100	197							
15:36:58	35	765	442	578	7	3	1	1	1			251	33	10	202	205	102	251	5	60	101	201	1	251	150	1			
15:36:59	35	748	433	578	5							246	33	10	198	201	98	246	5	59	98	198							
15:37:00	35	772	442	572	7	3	1	1	1			255	33	12	201	207	103	255	6	63	102	200	1	255	150	1			
15:37:01	35	753	439	582	5							245	33	10	200	204	99	245	5	59	99	200							
15:37:02	21	424	260	468	5							160	23	10	128	134	71	160	5	48	69	142							
15:37:03	11	189	136																										
15:37:04	11	193	142																										
15:37:05	10	177	125																										
15:37:06	8	145	102																										
15:37:07	7	133	83																										
15:37:08	4	102	54																										
15:37:09	4	100	53																										
15:37:10	3	91	40																										
15:37:11	3	96	36																										
15:37:12	3	94	36																										
15:37:13	4	99	48																										
15:37:14	4	100	53																										
15:37:15	4	100	53																										
15:37:16	3	81	40																										
15:37:17	4	310	55																										
15:37:18	17	1907	207																										
15:37:19	11	1034	134																										
15:37:20	4	100	53																										
15:37:21	8	270	100																										
15:37:22	11	385	133																										
15:37:23	11	387	133																										
15:37:24	11	267	135																										
15:37:25	11	186	136																										
15:37:26	11	194	140																										
15:37:27	6	135	80																										
15:37:28	2	78	27																										
15:37:29	2	76	29																										
15:37:30	2	82	30																										
15:37:31	2	84	31																										
15:37:32	8	147	94																										
15:37:33	11	188	136																										
15:37:34	11	186	136																										

Prepared by HP Customer
Prepared by Hewlett-Packard Co.
Customer Network Troubleshooting Center

« Throughput kbps »
Frame Relay Link Stats - 5 Minute Averages - DTE / DCE

□ Throughput ave (DTE) ■ Throughput ave (DCE)

HP Internet Reporter - ver 0.04.00 Release: 80PCEN0.RWS Chart Type: 31 Created: 82595 # 3:00:00PM

Contents

I. Introduction	3
II. Case studies	4
1. Documenting the state of the LAN	4
Mini-lesson: Sample periods and sample intervals	4
Mini-lesson: Using the Node Stats measurement to identify major nodes	5
2. Identifying power users and LAN abusers	8
3. Optimizing WAN application performance the scientific way	11
4. Identifying error rate trends with the HP Internet Reporter	16
5. Controlling costs and improving service on frame relay and X.25 networks	18
6. Verifying and documenting WAN circuit quality for trouble-free installations	22

I. Introduction

We live in a vastly reactive world. If it ain't broke, don't fix it. The widespread success of dispatched analyzers like Hewlett-Packard's Internet Advisor for LAN and WAN are a testament to this idea. Loaded with dozens of applications, the Internet Advisors provide their users with equally many ways to fix network problems as they occur.

Of course, this reactive approach results in too much downtime for today's network applications, which creates the need for proactive network management tools. HP's proactive solutions include OpenView, ProbeView, and NetMetrix. With such systems in place, network managers not only get alerted to problems, they can continuously monitor the operating characteristics of the network. With the proper understanding of these characteristics, network managers can improve network performance, lower operating costs, and reduce downtime.

The downside of proactive network management systems for many applications is the time and cost associated with installation. Although proactive systems are essential for managing large internetworks, cost-justifying their installation may be difficult for most managers of smaller networks, especially if all they need are periodic network health checks or performance test results.

With the addition of HP Internet Reporter, a software add-on, the HP Internet Advisor becomes a cost-effective alternative to permanently installed network management systems for performing network baseline studies. In this product pair, the Internet

Advisor collects statistical data then passes it on to the Internet Reporter for analysis. Internet Reporter produces a compact table of the data, adds interpretive statistics, then produces dozens of presentation-quality charts.

What is a network baseline?

A network baseline is a snapshot of various network statistics, trended or counted at specific intervals. Data is collected on the network over a representative period of time, from a few hours to several days. A basic baseline might detail only utilization and error rate statistics, while a more comprehensive baseline details dozens of statistics. Such statistics might include broadcast frames, management frames, or frames by specific protocol types, such as IP or SNA. Even more comprehensive baseline studies provide complete traffic breakouts of some or all stations on the network.

What is a network benchmark?

Whereas a baseline details the state of an entire network segment, a benchmark focuses attention to specific conversations. Benchmarks are typically run for just a few minutes to an hour, and at much shorter sample intervals than baselines. Any time you need to document or compare the performance characteristics of a networked device or application, a benchmark study will provide the facts you need.

Network baselines and benchmarks have a broad range of applications. We have detailed a few of the most common ones as case studies in this document.

Case Studies

Case 1: Documenting the state of the LAN

The starting point for understanding your network using Internet Reporter is with a general network baseline. This study will provide you with an invaluable reference document. In it, you will see a high-level overview of your network, complete with statistical trends and counts documenting utilization, errors, broadcasts, routing, protocols, stations, and more.

HP Internet Reporter solution

Three measurements on the HP Internet Advisor for LAN are capable of generating data files for the Internet Reporter. As you are trying to get an overview of

the network, it is important that you run these three measurements concurrently. Each collects statistical data in different ways.

The *Summary Stats* measurement collects trended data over time on a wide range of network interface information. This data includes utilization statistics, detailed network-specific error breakouts, and additional data such as token-ring source routing profiles and FDDI ring state information.

The *Protocol Stats* measurement collects both trended data and total counts over time, showing network traffic in terms of specific protocols. Counts even provide comprehensive frame size distributions. Protocol Stats can be configured to show statistics by protocol family, actual DLL stack type, IP protocols and ports, or Novell types and sockets.

The *Top Talkers* measurement provides counts for the 50 most active stations on an Ethernet or token-ring network.

Mini-lesson: sample intervals and sample periods

Sample intervals and sample periods are used in various combinations when trending Summary Stats or Protocol Stats data. It is imperative that you decide on the right values for your needs *before* you start collecting data.

Sample interval: The time between entries in a Summary Statistics or Protocol Statistics Trends file. A 1-second sample interval places an entry in the table every second. A 10-second sample interval places an entry in the table every 10 seconds.

Sample period: The total amount of time between the first and last entries in a Summary Statistics or Protocol Statistics Trends file.

Choosing a sample interval

Figure 1.1 provides a cross-reference detailing all supported sample interval / sample period combinations.

When baselining a network segment, you are generally more interested in obtaining an overview of network performance than specific per-second detail. If you were to chart 1-second samples over an 8-hour day, you would have 28,000 data points. The resulting Internet Reporter table would be hundreds of pages long, charts

would be unreadable, and analysis would be difficult, if not impossible.

A more practical sample interval of 1-minute would yield only 480 samples in an 8-hour day, or 1440 samples over a 24-hour period. At these intervals, your tables and charts process much faster (2 to 5 minutes), with better looking reports.

Sample Interval	Recommended Sample Period	Summary Stats	Protocol Stats
1 second	up to 1 hour	YES	NO
10 second	15 minutes to 8 hours	YES	up to 4 hours
20 second	30 minutes to 8 hours	NO	YES
30 second	1 hour to 12 hours	NO	YES
1 minute	2 hour to 2 days	YES	up to 1 day
2 minute	4 hours 2 days	NO	YES
3 minute	6 hours to 3 days	NO	YES
6 minute	12 hours to 6 days	NO	YES
10 minute	1 day to 2 weeks	YES	NO
15 minute	2 days to 2 weeks	NO	YES
30 minute	4 days to 1 month	NO	YES
1 hour or greater	1 week or greater	60 minute	up to 1 day

Figure 1.1 Sample interval/sample period cross-reference table.

Where, when and how long?

Before you begin, you have to answer a few key questions that will affect the data you collect.

Where should I place the Advisor?

The Internet Advisor only collects statistics based on frames it sees on the LAN. Because it can't see through bridges, routers or switches, correct placement is essential. If you have access to only one Internet Advisor, you may need to move it from segment to segment on successive days to collect all the data you want.

When should I collect data?

You should collect data on both a typical day on the network and on a heavy day, perhaps during end-of-month processing or your busy Monday. This way, you'll have a reference for both normal and stressed conditions on the network.

How long should I collect data for?

Typically, a 24-hour period per segment is adequate for a general baseline. You want to make certain that you don't miss such events as the early morning power-up and late night backup.

Step one: collect data with the Internet Advisor

With the Internet Advisor located on the first segment you want to baseline, load or create a node list with the friendly names of some or all stations. Friendly names will make it easier to identify stations than their standard 12-digit hexadecimal physical addresses.

Next, configure the Summary Stats measurement to trend up to seven stations for the day. Good stations to choose include servers, hosts, gateways, and routers.

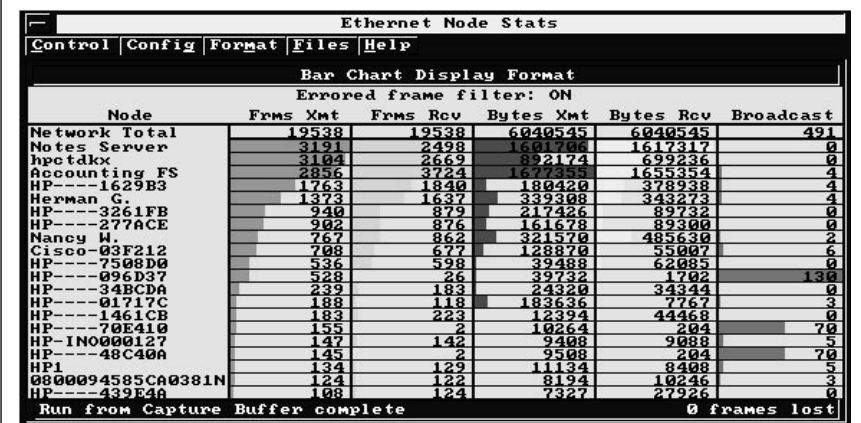
Mini-lesson: using the Node Stats measurement to identify major nodes

When trying to choose seven stations to specifically monitor in Summary Stats, you may have a difficult time identifying them by their 12-digit hexadecimal MAC addresses. Running Node Stats can really be a help here.

Without any filters active, run the Node Stats measurement, making sure that it is configured to sort on frames or bytes. Within just a few minutes, the major nodes on your network will be listed at the top of the table. (See figure 1.2.) This is the fastest way to get the MAC addresses of particular major nodes, especially if you are looking for the MAC address of a router.

If your network segment has too many major nodes, you may have to configure a filter to limit the ones you see at any given time. You can configure these filters from within the Node Stats measurement.

For instance, if you configure a filter to show only traffic to and from a particular user's workstation, you will see only those major nodes with which that workstation communicates.



The screenshot shows the 'Ethernet Node Stats' window with a menu bar (Control, Config, Format, Files, Help) and a toolbar (Bar Chart, Display, Format). Below the toolbar, it says 'Errored frame filter: ON'. The main area contains a table with the following columns: Node, Frms Xmt, Frms Rcv, Bytes Xmt, Bytes Rcv, and Broadcast. The data is sorted by Bytes Xmt in descending order.

Node	Frms Xmt	Frms Rcv	Bytes Xmt	Bytes Rcv	Broadcast
Network Total	19538	19538	6040545	6040545	491
Notes Server	3191	2498	1601708	1617317	0
hpctdkx	3104	2669	892174	699236	0
Accounting FS	2856	3724	1677488	1655354	4
HP---1629B3	1763	1840	189420	378938	4
Heyman G.	1373	1637	339308	343273	4
HP---3261FB	940	879	217426	89732	0
HP---277ACE	902	876	161678	89300	0
Nancy W.	767	862	321570	485630	2
Cisco-03F212	708	677	128870	35007	6
HP---7908D0	536	598	39488	6205	0
HP---096D37	528	26	39732	192	130
HP---34BCDA	239	183	24320	34344	0
HP---01717C	188	118	183636	7767	3
HP---1461CB	183	223	12394	44468	0
HP---70E410	155	2	19264	204	70
HP-IN0000127	147	142	9408	9088	5
HP---48C40A	145	2	9508	204	70
HP1	134	129	11134	8408	5
0800094585CA0381N	124	122	8194	10246	3
HP---439E40	108	124	7327	27926	0

At the bottom of the window, it says 'Run from Capture Buffer complete' and '0 frames lost'.

Figure 1.2 The HP Internet Advisor for LAN's Node Stats measurement helps identify major nodes on the network.

Case Studies

Remember, because the Internet Advisor collects statistics based on physical addresses, all devices on the far side of a router will be grouped together in a statistic for the physical address of the router port connected to the segment the Internet Advisor is on. Bridges, unlike routers, do not modify the physical layer source address field, allowing the Internet Advisor to “see” remote stations individually by their true addresses.

Be sure to configure both Summary Stats and Protocol Stats to log data to disk at 1-minute sample intervals. This interval will produce the most readable 24-hour baseline charts.

Make certain Protocol Stats is configured to *Show statistics for Data Link Layer* and that the *Show protocol stacks* configuration button is set. Also make sure that both Summary Stats and Top Talkers are configured to display stations by either their friendly names or their 12-digit hexadecimal addresses, depending on your preference.

Now, start running Summary Stats, Protocol Stats, and Top Talkers. You will want to begin either first thing in the morning, before your users arrive, or right before you go home. Ideally, starting at 12:00 midnight makes the best looking charts, although this scenario may be difficult to coordinate.

After they’ve run for 24 hours, stop the measurements. At this point, your logged trends files have already been saved to disk. Now, create a Protocol Stats total counts file (for a frame size

distribution) and a Top Talkers total counts file for use with Internet Reporter.

Repeat this entire step for each segment you need to baseline.

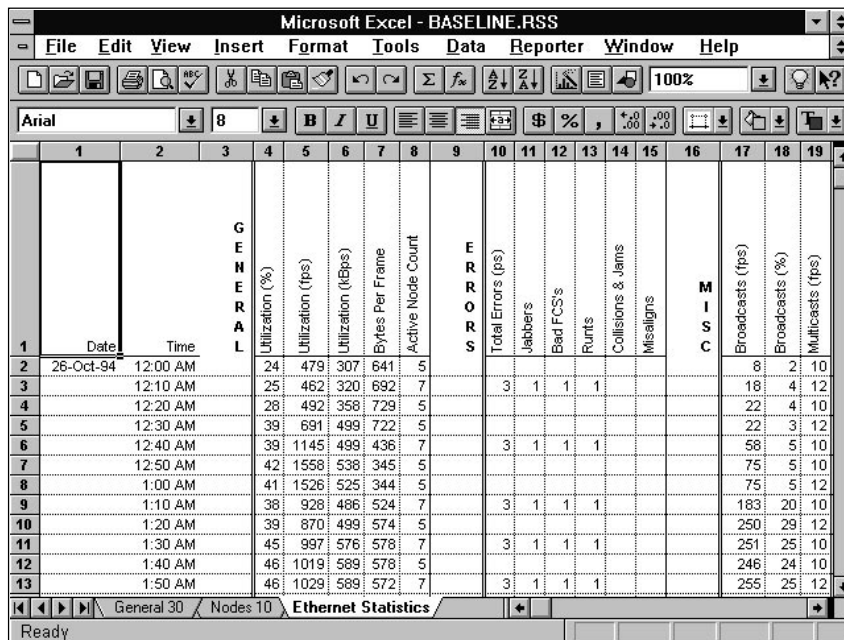


Figure 1.3 Data collected with the HP Internet Advisor is automatically reformatted by Internet Reporter's Table Builder into an easy-to-read spreadsheet.

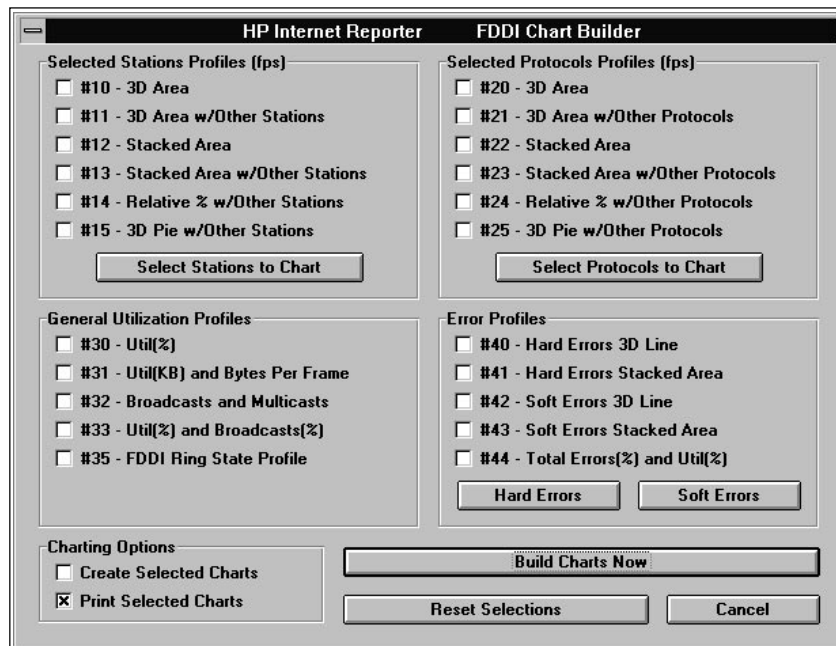


Figure 1.4 Internet Reporter's Chart Builder provides a point-and-click interface for creating a wide range of custom charts.

Step two: build tables and charts with Internet Reporter

With all your data files copied to a PC running Internet Reporter, run the Table Builder and Chart Builder applications to create a binder of reference documents for each segment. (See figures 1.3 and 1.4.) When you run Chart Builder, there is usually no need to display the charts on-screen. By selecting the "output to printer only" option, you will save time and disk space. Just be sure your printer is on-line and that any local spoolers are turned off.

Interpreting the results

By now, you have printed dozens of charts documenting all aspects of your network's traffic. (See figures 1.5 and 1.6.) Although each chart is clearly labeled and organized, you may not be sure what all the statistics mean.

Supplemental information about each statistic can be found in glossaries contained in the


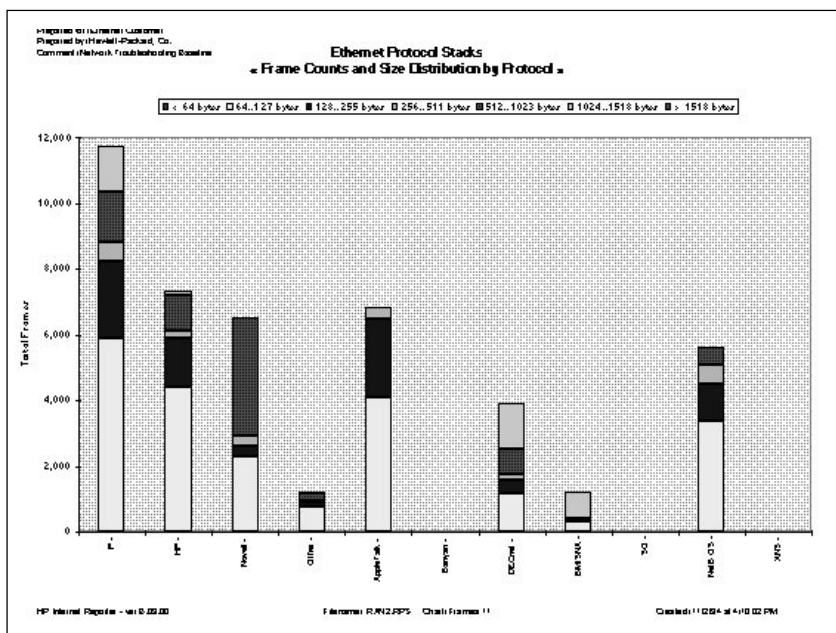


Figure 1.5 One of dozens of Summary Stats charts, this one shows network utilization and stations on the network trended over time.

HP Internet Reporter User's Guide. The glossaries provide definitions of all statistics and additional information useful in interpreting and troubleshooting specific conditions.

The documentation you now have can be used for many purposes, including cost-justifying network upgrades, use as a reference when problems arise, and for network segmentation.



Running baselines at regular intervals, say once a month, will show you general trends in the use of your network over time. The cumulative effect of these baselines will take much of the guesswork out of network planning.

Figure 1.6 This Protocol Stats chart provides both total frame counts and a frame size distribution for each major protocol family.

Case Studies

Case 2: Identifying power users and LAN abusers

In most networks, traffic is not evenly distributed across all user workstations. Different work habits, job descriptions, and workstation configurations will affect how much load each user places on the network. Understanding how each user utilizes the network helps you plan for growth and reconfigure for optimum performance. In the process, you may even uncover some LAN abusers, i.e., workstations consuming network resources at an unusually high rate.

Internet Reporter solution

The HP Internet Advisor for LAN's Top Talkers measurement will produce the bulk of the data required for this study. Top Talkers will count the total number of frames and Kbytes transmitted and received from each of the 50 most active stations on the network. After the top talkers are identified, specific suspect stations can be trended with the Advisor's Summary Stats measurement, providing further detail about those stations' network utilization.

Step one: collect data with the Internet Advisor

With the Internet Advisor located on the segment from which you want to collect data, run the Top Talkers measurement. The only configuration option you need to decide on before running the measurement is whether you want to track each station by its friendly name or its 12-digit hexadecimal address.

While friendly names are generally easier to identify with when

reading a table or chart, they do require that your node list be up-to-date and accurate. If you rely on the Node/Station Discovery measurement for friendly names, you may still not get the results you want. On a Novell network, for example, if many users share login names, you would not be able to differentiate them from each other. Hexadecimal addresses may be more difficult to read, but they can always be counted on for accuracy.

If your measurement is configured for friendly names, all stations not entered in your node/station list will be identified with a NIC vendor code (when available) plus the remaining unique 6-digit station identifier.

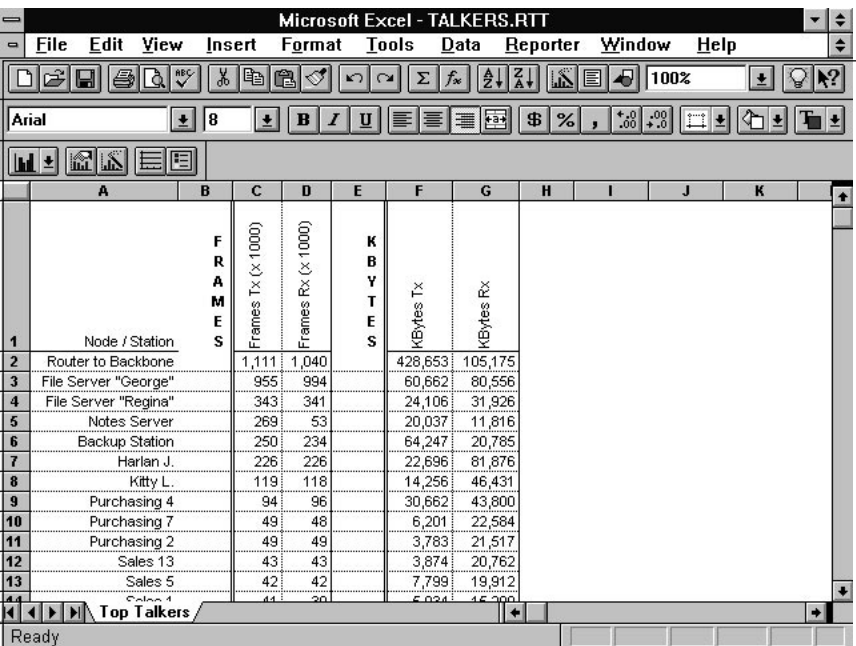
Now, run the Top Talkers measurement for a representative period of time. A single business day is adequate for a single file.

Running Top Talkers each day for a period of several days will make a difference in identifying "real" top talkers, and not just those stations which happened to have a single busy day. Because this measurement produces a single counts value for each statistic, creating a separate file for each day will prove useful.

Step two: build tables and charts with Internet Reporter

With all your data files copied to a PC running Internet Reporter, run the Table Builder and Chart Builder applications to create a set of tables and charts. (See figure 2.1.)

On all but peer-to-peer networks, you will notice that the most active top talkers are your major nodes. These include servers, hosts, gateways and routers. When you think about it, most network traffic is either sourced



The screenshot shows a Microsoft Excel spreadsheet titled "Microsoft Excel - TALKERS.RTT". The spreadsheet contains a table with the following data:

	A	B	C	D	E	F	G	H	I	J	K
	Node / Station	FRAMES	Frames Tx (x 1000)	Frames Rx (x 1000)	KBYTES	KBytes Tx	KBytes Rx				
1	Router to Backbone		1,111	1,040		428,653	105,175				
3	File Server "George"		955	994		60,662	80,556				
4	File Server "Regina"		343	341		24,106	31,926				
5	Notes Server		269	53		20,037	11,816				
6	Backup Station		250	234		64,247	20,785				
7	Harlan J.		226	226		22,696	81,876				
8	Kitty L.		119	118		14,256	46,431				
9	Purchasing 4		94	96		30,662	43,800				
10	Purchasing 7		49	48		6,201	22,584				
11	Purchasing 2		49	49		3,783	21,517				
12	Sales 13		43	43		3,874	20,762				
13	Sales 5		42	42		7,799	19,912				
44	Sales 4		41	20		6,024	15,000				

Figure 2.1 Internet Reporter builds an organized table detailing the activity of your most active 50 stations.

from or destined to one of these devices. While important for a baseline study, these major nodes are not significant for our purposes here. Instead, their high values will tend to flatten the workstations' counts on charts, making them very difficult to read.

To solve this problem, you should delete these major nodes from the Internet Reporter table by highlighting their individual row numbers and selecting *Delete* from the Excel *Edit* menu.

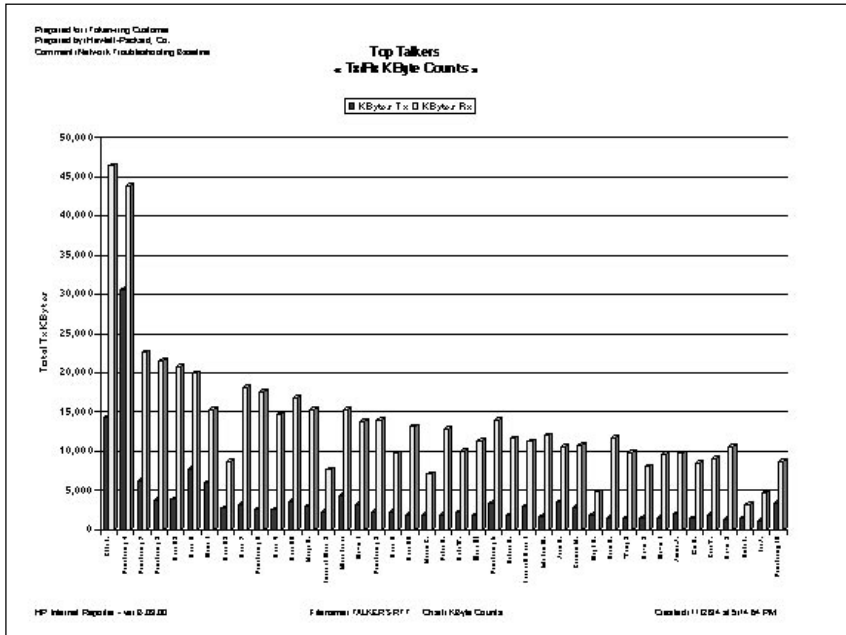


Figure 2.2 Top Talkers counts charts help you quickly identify the power users and LAN abusers on your network.

Interpreting the results

When looking at one of the Total Counts charts, you will see individual station counts by either frames or Kbytes. (See figure 2.2.) It is important to differentiate between frame and Kbyte utilization.

Different applications and protocols generate different size frames. A station showing up as a top talker in terms of frames may not generate much in the way of Kbytes at all if it is running primarily VT-100 terminal emulation, which is generally carried in lots of very small frames. Likewise, a station on a token-ring network could be moving large blocks of data with an 8 Kbyte frame size, strongly taxing the network. Because of the large frame size, however, this station will only show up as a top talker in terms of Kbytes, and not frames.

Analyzing the charts is a matter of asking yourself questions and searching for answers. Ask yourself the following:

- Q.** Do workstations with common functions exhibit similar utilization?
 - A.** If, for example, eight order entry workstations are grouped together, but another two have very different utilization counts, you need to find out why. Users grouped outside of their peers in terms of network utilization should be marked as suspect, then have the following question researched.
- Q.** What is causing a particular station to generate an abnormally high percentage of the traffic?
 - A.** The reasons vary from network to network. In a Windows environment, a swap file could be incorrectly configured to run on a file server. Perhaps the user runs different applications or was running reports that day. Maybe the workstation has a faster or slower NIC or CPU. Or possibly the user is just much more, or less, productive than the others.

The possibilities are endless. Internet Reporter will identify *which* stations are a problem, but you still have to research *why*. Using the HP Internet Advisor's other decode, Commentator, and statistics measurements, along with traditional troubleshooting methods, can help you pinpoint the real causes.

Read/write ratio charts

In addition to identifying the top talkers, Internet Reporter lets you see whether a particular station is primarily a read or write device, based on the ratio of transmit to receive frames and Kbytes. (See figure 2.3.) This information is very important in determining if a station is configured properly.

Let's say a particular station is generating much more traffic than its peers performing similar functions. Looking at the ratio of transmit to receive frames/Kbyte might provide more insight. A station with a much higher

Case Studies

percentage of receive frames may indicate that it is configured to load all of its applications from a server, while the other stations load applications from local hard disks.

Step three: trending suspect workstations

Based on the findings in step two, you should choose seven stations to be specifically trended over time to uncover their traffic patterns over the course of a day. Pick a combination of suspect stations and “normal” stations with which to compare them.

Then, run Summary Stats and Top Talkers with the new configuration for a day and study the results. For this study, a 10-minute sample interval for 24 hours should be adequate, as you are just looking for overall trends for these seven stations and do not need to capture activity spikes.

Interpreting the results

Looking at the trends of the top talkers identified in step two shows some interesting results. (See figure 2.4.)

Looking again at an unusually talkative station might reveal that it performs very much like its peers until late in the day, between 4 and 8 p.m. Perhaps this station runs batch jobs or other network-intensive processes during these hours. Further investigation is necessary at this point, but at least Internet Reporter has pointed you in the right direction. If it turns out that this user’s jobs can be run later at night, you can return valuable bandwidth to other users at the end of the day.

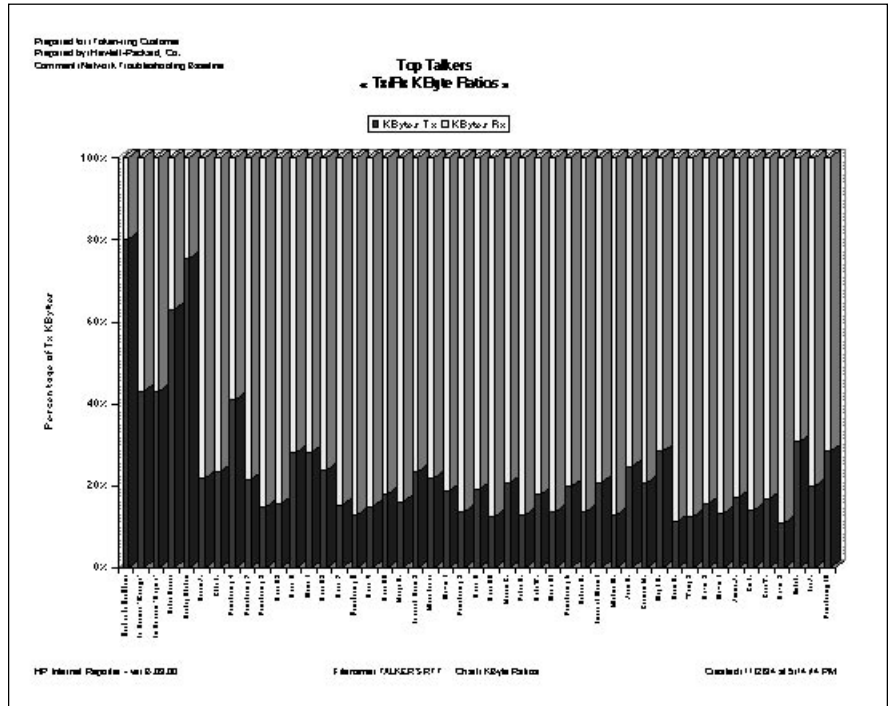


Figure 2.3 Internet Reporter lets you graphically see the ratio of received data to transmitted data by station in terms of frames or Kbytes.

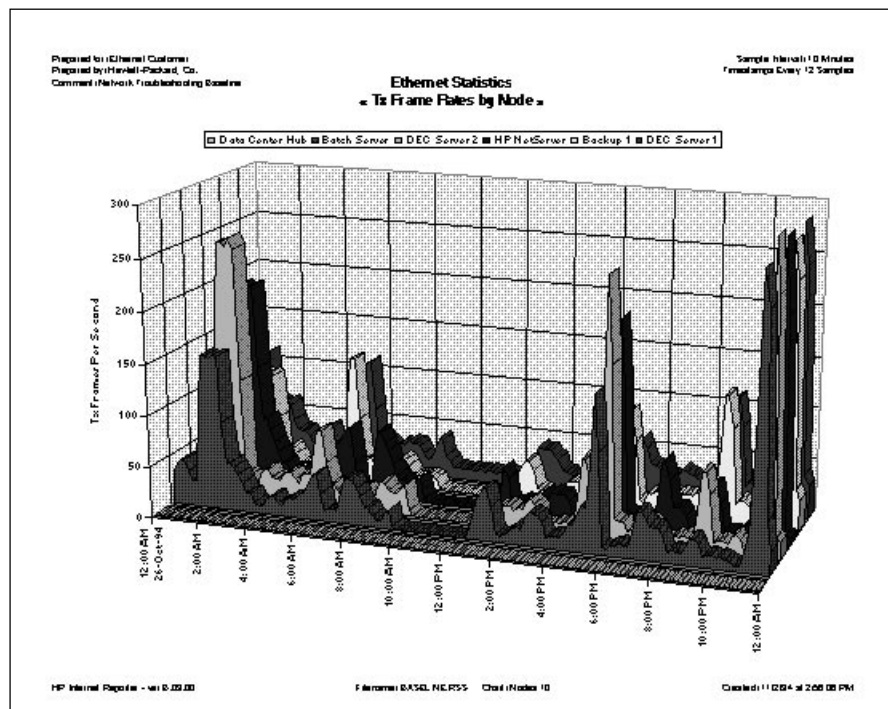


Figure 2.4 Summary Statistics trends charts characterize station traffic based on activity at different times of the day.

Case 3: Optimizing WAN application performance the scientific way

You've just been told that 80 remote offices need to access the new corporate imaging system. At first, you cringe at the prospect of squeezing pictures through telephone wire. But when it comes time to figure out a solution, Internet Reporter proves to be a life saver.

Many issues must be dealt with. Some you can control and others you can't. In this case, we assume that you have complete flexibility in WAN circuit configurations but are bound to a specific application for delivering images to the desktop.

Internet Reporter solution
To solve this problem, you will use the HP Internet Advisor for LAN and HP Internet Reporter in a test lab environment as a benchmarking tool, documenting the results of a variety of testing scenarios. Comparing the results of these tests will help you decide on the best solution for the problem.

Although you could perform the tests with an HP Internet Advisor for WAN, you would not be able to differentiate between the different test workstations as easily. On the LAN, each workstation is identified by its own network address; on the WAN, each workstation is identified by the router's address.

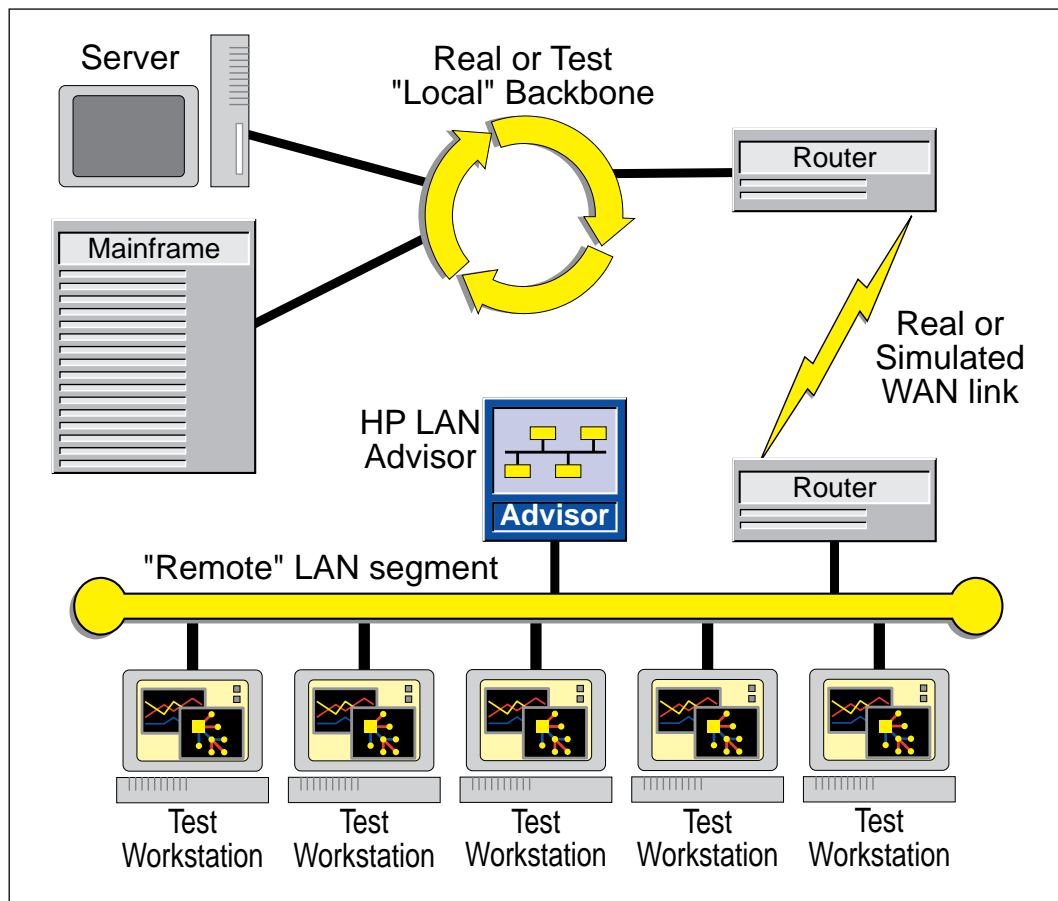


Figure 3.1 Sample WAN application benchmarking lab diagram.

Case Studies

Step one: setting up the test environment

To identify the best solution, you need to sample as many different configurations as possible. This means setting up a lab environment that you can quickly reconfigure between tests, without disrupting (or being affected by) your regular network traffic.

At a minimum, you should connect a couple of routers to each other with a test cable capable of simulating any speed line. (See figure 3.1.) For testing, you will also need at least one representative workstation on the “remote” segment. Having access to several remote stations is best, so that you can perform multi-station tests. If you can’t reproduce the actual host devices in the lab, you will need to connect the “local” router directly to the backbone segment on which the real hosts are located.

Step two: composing a game plan

During the testing process, you will potentially perform dozens of tests, each with an array of tables and charts created from them. Without a game plan, it will be impossible to keep track of your data, let alone guarantee that you tested every combination you intended.

The combination of tests you run will be based on the information you receive during a fact-finding process. This process centers around interviewing key people involved in the project, including the end users, department managers, MIS staff, consultants, and vendors.

In this example, after talking to the imaging-application vendor, you find that, although many of the front-end applications can be loaded from either the host or the workstation, the vendor is not sure how much performance you will yield by loading the applications at the remote site. The vendor also tells you that you have a choice of accessing the images on the host with either direct emulation or through a gateway, using a variety of protocols.

Talking with a few of the branch office managers brings up the issue of workstation performance. Locally connected Pentium™ -class machines outperform the 486-class ones, but no one knows if there will be a significant, or any, difference on the WAN. The managers also have questions about workstation RAM and brands of network interface cards.

With so many variables and so many tests to choose from, you have to narrow down your test selections to a manageable lot, while still accomplishing as many test scenarios as possible. A worksheet like the one in figure 3.2 will provide the organization required for this project.

Because of the obvious WAN performance benefits of loading applications locally at the remote sites, all tests will be performed this way. Also, although there is a choice of protocols from the application vendor, internal standards dictate the use of TCP/IP to get to the host.

Step three: collecting the data

With a game plan in hand, it is time to set up the environment and begin collecting data with the HP Internet Advisor. The Summary Stats measurement will be used to collect data during each test, as it is capable of trending data for specific stations over time. Notice the two basic types of tests in the sample Test Summary Worksheet. The first test is very controlled and has a short duration. This will allow you to run it in all the different ways listed in a short period of time, with all other variables removed. The second test simulates, as best as can be expected in a lab, normal usage of the system in a variety of ways. As these tests have a longer duration and require real users to visit the lab, there is only time to complete a few variations.

Be certain to connect the Internet Advisor to the “remote” segment, so that you do not have to configure filters to limit the scope of your tests to the lab. Also be sure to manually configure Summary Stats for the stations you will be testing. If you use the *Discover new nodes/stations* option, your files will not be compatible with Internet Reporter. If, at any time, you stop Summary Stats and restart it without creating a new log file, your original file will be appended. To prevent this, make a habit of creating a new log file immediately before each test.

HP Internet Reporter Benchmark Test Summary Worksheet

Test #	Test Summary	# of Stations	Line Speed	Host Via	PC Type	RAM Amount	NIC Type
1	A: Load application B: Select image browser button C: Search range of documents 100 to 3500 D: Request image document between 180 and 190 (each station chooses unique #, all are same size) Sample Interval = 1 second Sample Period = less than 3 minutes						
1.1.1		1	56KB	Direct	486	16MB	16-bit
1.1.2		2	56KB	Direct	486	16MB	16-bit
1.1.3		5	56KB	Direct	486	16MB	16-bit
1.2.1		1	56KB	Direct	Pentium	16MB	16-bit
1.2.2		2	56KB	Direct	Pentium	16MB	16-bit
1.2.3		5	56KB	Direct	Pentium	16MB	16-bit
1.3.1		1	256KB	Direct	486	16MB	16-bit
1.3.2		2	256KB	Direct	486	16MB	16-bit
1.3.3		5	256KB	Direct	486	16MB	16-bit
1.3.4		5	T1	Direct	486	16MB	16-bit
1.4.1		1	256KB	Direct	486	16MB	32-bit
1.4.2		2	256KB	Direct	486	16MB	32-bit
1.4.3		5	256KB	Direct	486	16MB	32-bit
1.4.4		5	T1	Direct	486	16MB	32-bit
1.5.1		1	256KB	Direct	486	16MB	16-bit
1.5.2		1	256KB	Direct	Pentium	16MB	32-bit
1.5.3		1	256KB	Direct	Pentium	16MB	32-bit
1.5.4		1	256KB	Gateway	486	16MB	16-bit
1.5.5		1	256KB	Gateway	Pentium	16MB	16-bit
1.5.6		1	256KB	Gateway	Pentium	16MB	32-bit
2	Random I/O with test users sitting in Lab Sample Interval = 1 second Sample Period = 15 minutes						
2.1		5	56KB	Direct	Pentium	16MB	16-bit
2.2		5	256KB	Direct	486	16MB	16-bit
2.3		5	T1	Direct	486	16MB	16-bit
2.4		5	56KB	Direct	Pentium	16MB	32-bit
2.5		5	256KB	Direct	Pentium	16MB	32-bit
2.6		5	T1	Direct	486	16MB	32-bit

Figure 3.2 Every successful benchmark starts with a game plan and a Test Summary Worksheet.

Case Studies

Test one: 4-step short test

For this series of 20 short tests, one data file will be created for each. To visually separate the four steps in each test file, you will have to pause for 15 seconds between each step. This quiet time will manifest itself as a flat line in your charts. (See figure 3.3.)

If you want to condense the number of files required for this test, you can run several of them back to back by configuring each of the seven Summary Stats stations to a unique station on the remote segment. You will still have to wait that 15-second period between tests and steps, but you can reduce the number of files and charts you have to review. Figure 3.4 illustrates the results of back-to-back testing in a single file.

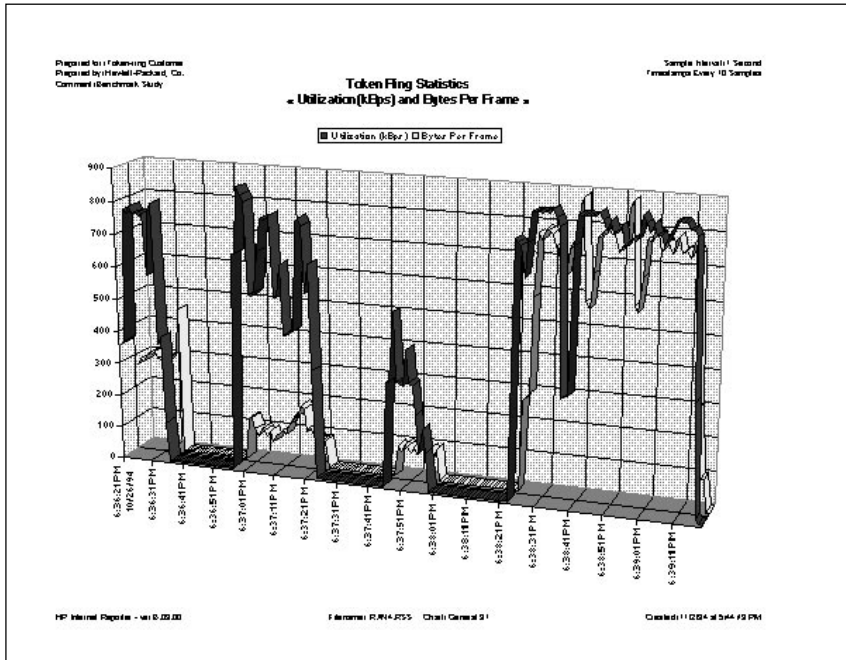


Figure 3.3 This utilization chart clearly illustrates the results of each of the four steps in test 1 on the Test Summary Worksheet.

Test two: random i/o test

The most difficult part of this test is getting a group of users to volunteer an hour or two of their time to simulate or perform their work day in your lab. While the first set of tests documented potential system capabilities, the information derived from these tests will give you an understanding of how the system will hold up in a real situation.

Because of the number of samples collected by each of these tests, you will not want to group multiple tests together in files. Instead, keep a 1-to-1 ratio of files to tests.

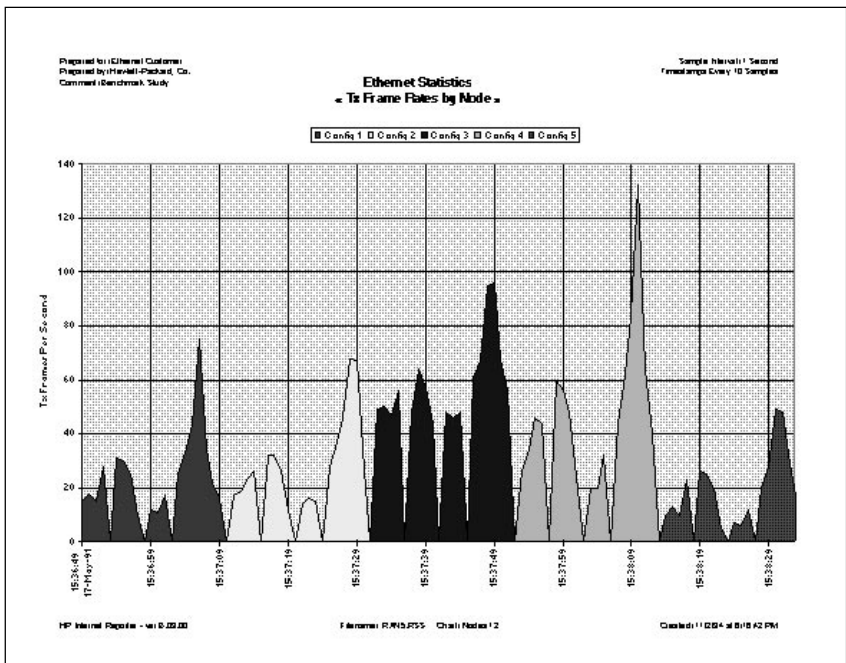


Figure 3.4 By running several different tests back-to-back in the same log file, you can see the results together in the same chart.

Step four: processing the data

This is the easy part. Run each of the data files through Internet Reporter's Table Builder and Chart Builder to produce documentation of your test results. When creating your documentation, be sure to use Internet Reporter's comment fields to describe each test.

Be certain to keep a printed copy of each master table with its related charts, as the tables will provide a reference of exact values that may not be discernible on certain charts.

Interpreting the results

Using your test summary worksheet as a reference, compare the results of the various tests. In the charts will be all of the tradeoffs between your different configurations.

Pay particular attention to the correlation of line speeds to the number and type of stations in the test. You will notice that with slower speed lines, you will reach capacity much more quickly than with faster lines. So a Pentium machine that outperforms a 486 in a single station test may not at 5 stations. This is due to the limiting factor of the slow line.

On the other hand, your real-world random i/o test may show that users have very sporadic activity, each using the link at different times. With this knowledge, you can determine that in your environment, investing in Pentium machines will indeed be beneficial.

Conclusion

The bottom line in benchmarking is to pay attention to the numbers and to read between the lines. Doing so, you can make decisions with much more confidence than you had before using HP Internet Reporter.

Case Studies

Case 4: Identifying error rate trends with the Internet Reporter

Although not normally considered a troubleshooting tool, Internet Reporter's exceptional trending of error statistics makes it ideally suited to identifying long-term network problems.

All Ethernet errors are grouped together in Internet Reporter tables and charts, while token-ring and FDDI errors are grouped into hard-error and soft-error categories.

Both token-ring and FDDI have a management layer that assists in the diagnostic process. Ethernet lacks this management layer and can be very difficult to troubleshoot as a result. This case study focuses on understanding more about Ethernet errors and how to use Internet Reporter to interpret them.

Reading the charts

When the HP Internet Advisor for LAN collects Ethernet error statistics, it categorizes them by type, and provides a total count of errors. The six Ethernet error statistics tracked by the HP Internet Advisor and Internet Reporter are as follows:

Collisions and Jams

These are the most well-known and most misunderstood of Ethernet errors.

Collisions are a regular Ethernet event and occur when two or more nodes try to send on the Ethernet simultaneously. When a collision occurs, the sending

nodes sense the collision, perform a random count, and retry. As the network becomes more heavily loaded with frames, more collisions will occur. High collision rates can also be caused by faulty adapters or nodes generating frames out of control and saturating a network.

Jam signals are sent by transmitting nodes when they sense a collision to keep other nodes from transmitting. On a 10Base-T hub, a collision on any one port will cause a jam signal to be sent out on all hubs' ports. For a detailed breakout of network collisions and jams, run the Ethernet Vital Signs measurement at the same time and duration as the Summary Stats measurement.

Total Errors (ps) - *Total errors represented in errors per second.*

Total Errors is a total of the following statistics: runts, jammers, Bad FCSs, and misaligns.

Notice that collisions and jams are not counted here because, technically speaking, they are not errored *frames*; they are simply error *conditions* on the network.

Also note that an individual frame may have several types of errors. For example, if an individual frame is a runt, and has a bad FCS, it will count as two errors (one for the runt and one for the bad FCS). If the errored frame was the result of a collision, the collisions and jams statistic will also have a count in it.

Runts

Runts are frames that are shorter than 64 bytes and therefore invalid on an Ethernet network. They can be the result of collisions on the network, or they can be a sign that a node is generating short frames without padding them up to 64 bytes. Because runts usually are too short to include a source address, they are very difficult to associate with a particular node.

Jammers

Jammers are frames that are longer than 1518 bytes, and therefore invalid on an Ethernet network. Jammers are usually the result of either a node generating frames outside Ethernet specifications or a faulty transceiver on the network.

Bad FCSs

These are frames with a frame check sequence that does not pass a checksum calculated when the frame was transmitted. This is usually an indication of a faulty transceiver or cable system component. If so much as one bit is modified during transmission by equipment or electrical problems on the Ethernet, the frame will have an FCS error. Collisions will cause frames to have a bad FCS as well.

Misaligns

Misaligns are frames whose length in bits is not divisible by eight (as in eight bits per octet). These frames also have bad FCSs and are usually the result of electrical problems on the network cabling, a faulty workstation, or the result of a collision.

Example

Looking at the chart in figure 4.1, you can see that this network has a significant number of collisions and jams but does not have other Ethernet errors in correlating numbers. This is a sign of a very busy network, as high collision

rates typically indicate high utilization. However, because the other error types are low, we can conclude the network does not have other problems affecting performance, such as wiring faults or bad NICs.

Analyzing the relationship of different error types to each other will help in identifying the real causes of your Ethernet errors.

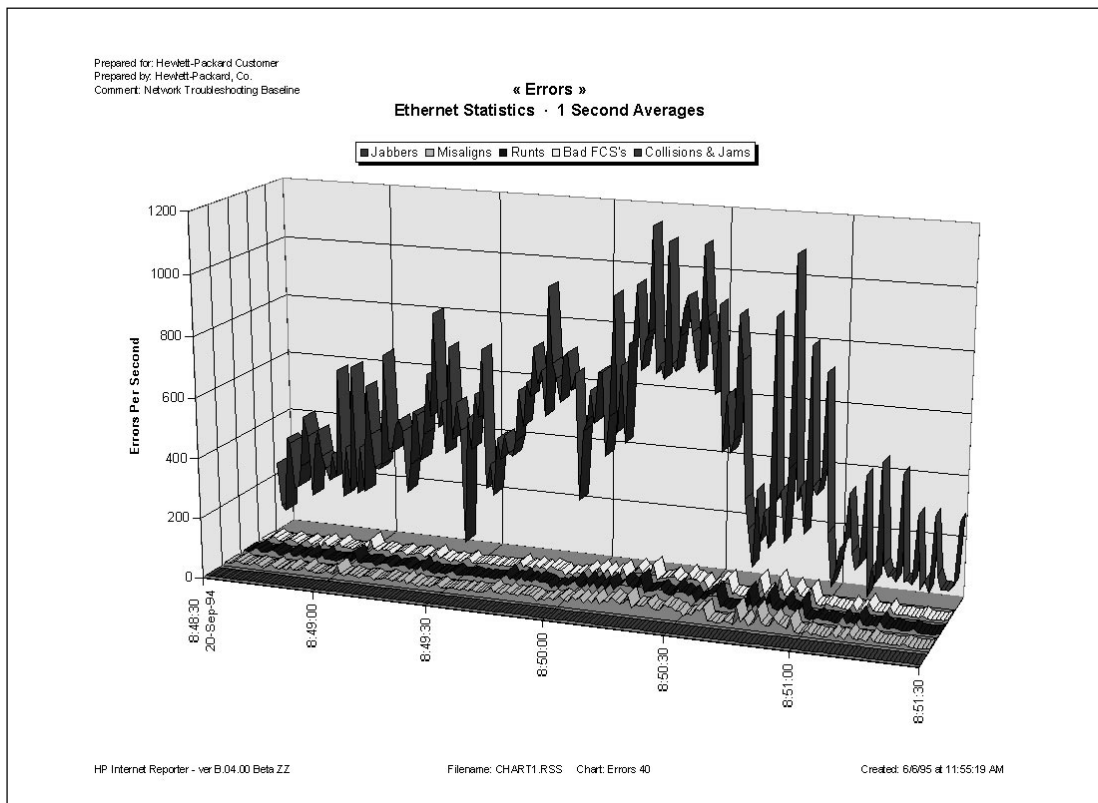


Figure 4.1 On this network, the high collision rate indicates a saturated network. However, because the other error statistics are low, we can conclude that the network is otherwise healthy.

Case Studies

Case 5: Controlling costs and improving service on frame relay and X.25 networks

X.25, frame relay, and other new virtual-circuit based networks have dramatically improved the performance and cost-effectiveness of providing multiple-site WAN services to customers. Along with the flexibility “the cloud” affords comes a host of difficult bandwidth management problems. Just ask yourself the following questions and you’ll get the idea.

- Are my virtual circuits configured for optimum throughput?
- Is my service provider delivering the information rates I’m paying for?
- Is there traffic on my network that shouldn’t be there?
- How much additional traffic can my network handle?

Most network clouds are designed by either adapting older point-to-point designs, educated guessing, or a combination of both. These are relatively low-tech methods for high-tech problems. And unlike the good old mainframe days, network resources are often distributed across the cloud instead of in a central data center, making network design an even greater challenge.

Internet Reporter solution

In case study 3, “Optimizing WAN application performance the scientific way,” you saw ways to better estimate bandwidth requirements for your network design by using Internet Reporter to benchmark application performance. Now that the cloud is in place, you can use Internet Reporter to continually fine tune your network to get the most throughput for your dollar.

You’ll start by collecting statistical data at critical points around the network cloud. For this example, a frame relay network with 10 physical circuits and 11 virtual circuits will be used. Then you’ll use Internet Reporter to plot and analyze the data. From these charts, you can make recommendations for changes to the network, and you may even identify other network-related problems.

Step one: choosing physical circuits to baseline

If your network has just a few physical circuits, you will probably want to baseline each one. However, on this network, (see figure 5.1) you will pick just a few critical connections to baseline. Of the 10 physical connections to the cloud, four support most of the users and applications, while the other six support distribution warehouses with just a few terminals each. Further examination of the network shows that all of the

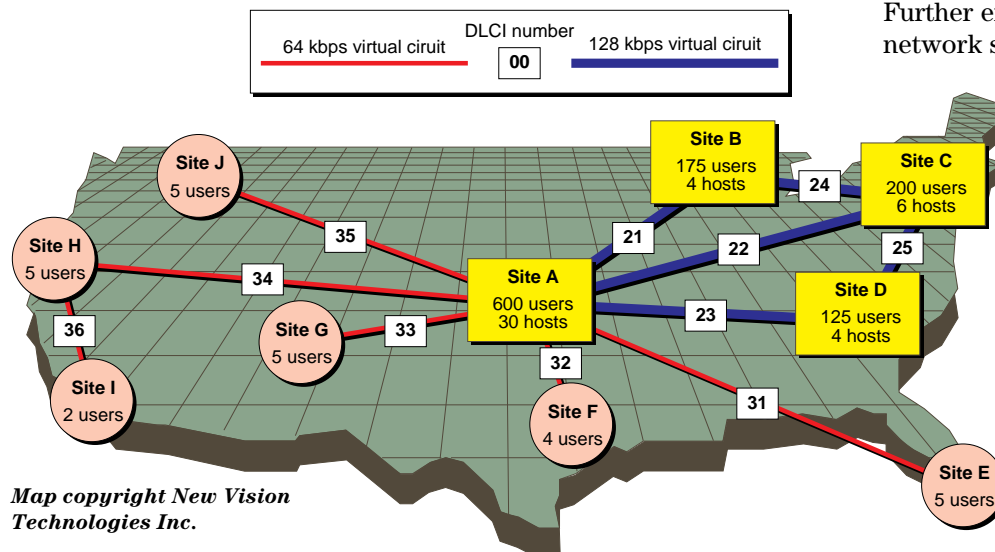


Figure 5.1 This frame relay network utilizes a mesh of 128 kbps virtual circuits (DLCIs) between the four main sites, with slower 64 kbps virtual circuits to six warehouses.

significant virtual circuits have at least one termination at either site A or C. Because Internet Reporter Can break-out individual virtual circuits, you can baseline at just two physical sites (A and C) to collect statistical data on all of the critical virtual circuits. A more comprehensive study would include the four main sites A through D.

Step two: collect data with the Internet Advisor

Before you begin collecting data, you must first make a few decisions about the installation and configuration of the Internet Advisors.

First, you must decide on a time to collect statistical data. Most baselines of this type are run for 24 hours, allowing enough time to profile a full day's processing. Depending on your needs, you may need to collect data for several days. If you do, try to break up each day into separate log files so that you can compare each day side-by-side when you chart them.

In a perfect world, you would collect data at all sites concurrently to guarantee statistical accuracy. However, in the real world, you may have access to only one or two Internet Advisors, with several physical circuits to baseline. In these cases, you should schedule the sites on different days that you are confident have similar traffic patterns and loads. A bad idea would be to baseline one circuit during year-end processing and another circuit just before a holiday.

Second, you need to decide where in the chain of network equipment to connect each Internet Advisor. Try to locate the Internet Advisor at a point closest to the network cloud. Depending on the premises equipment you have installed, you may have to disconnect the circuit in order to insert the Internet Advisor. Some equipment, such as managed CSUs, have monitor jacks that you can connect to without disrupting service.

In either case, be aware that your service provider may not allow you to connect directly to their equipment or lines without specific permission! If it is not practical, or permissible, to connect this close to the network cloud, you will have to connect through one of the V-series interfaces between such devices as a router and a DSU. Plan on bringing down the connection for the time it takes to connect the Internet Advisor into the V-series cabling (about one minute). Also, be aware that connecting through a V-series interface will not allow you to collect statistics on T1 or CEPT/E1 specific error conditions, but will still provide all other statistics for Internet Reporter.

Third, you should choose a program from the Internet Advisor Toolkit that will provide additional counters that may be useful for a baseline. Frame-relay specific programs included with the Toolkit include five programs that are based on different frame relay specifications, one for LAN over WAN, and one for TCP/IP

over WAN. This example deals with basic utilization and throughput only, so the counters you choose are not important.

Last, you must decide on a sample interval and sample period for your baselines. These are also known as the logging interval and logging period in most WAN programs. The Internet Advisor allows you to collect statistical data at regular intervals for a specific length of time. Although the Internet Advisor will let you select sample intervals in one-minute increments up to several hours, usually one to ten minutes will be appropriate. For this example, you will collect data at five-minute intervals for a period of 24 hours. This will yield 288 samples, making charts easy to read and analyze.

Case Studies

When choosing a sample interval or sample period combination, be careful not to create a log file that will have too many samples. If your log file contains more than a few hundred samples, report processing times grow long and charts become increasingly more difficult to read. Because the Internet Advisor collects minimum, average, and maximum utilization and throughput values for each sample, you do not have to worry about missing the peaks and valleys present in the statistics when using a long sample interval. Most other WAN statistics such as errors and counters are collected as total counts within the sample interval instead of being reduced to average per-second values as in Internet Advisor for LAN software.

With these preparations completed, you can begin collecting statistical data according to your deployment schedule. At the end

of each 24-hour period, you will copy the log files to a computer running Internet Reporter and free up the Internet Advisor for other tasks.

Step three: plot the results with Internet Reporter

With the log files on the hard disk of a computer running Internet Reporter, you will run Internet Reporter's Table Builder to build a series of data tables that will, in turn, be used to create charts. One table will be created as a summary of the entire physical circuit (link). Additional tables will be created for each of the virtual circuits (DLCIs). A final table will be created as a cross-reference of the various virtual circuits.

During the process of building the tables, you will be prompted to input the Committed Information Rate (CIR) for each DLCI on

the network. (See figure 5.2.)

These values are never transmitted on the actual circuits, so you must input the values that you have contracted for with your service provider. Later, you can go back and modify these CIRs to see their effect on the network cloud.

Step four: interpreting the results

The first chart you will review is one that shows utilization of the total physical link. (See figure 5.3.) This baseline shows that site A is at a busy but healthy level, averaging 70 percent utilization during peak hours, with maximum values taking full advantage of the line. Site C, not shown here, is operating at about the same levels.

While the overall link appears healthy, a review of the individual DLCI charts shows a different picture. Most of the virtual circuits appear to be operating within acceptable ranges, with the exception of DLCI 23, which connects sites A and D. (See figure 5.4.) This chart shows that the CIR for the circuit is much too low for the level of traffic trying to get across it. There are long periods of sustained traffic near 100 percent, and many bursts well over that. Looking at charts for other DLCIs not shown, we also notice that the 128 kbps CIR between sites B and C is barely utilized, but should be left alone as it is an important alternate path within the "mesh" of the four main sites.

DLCI #	CIR (kbps)	DLCI #	CIR (kbps)	DLCI #	CIR (kbps)
21	128				
22	128				
23	128				
31	64				
32	64				
33	64				
34	64				
35	64				

Figure 5.2 Internet Reporter lets you enter the Committed Information Rate (CIR) for each virtual circuit.

In this simplified example, you have discovered that one of the virtual circuits, DLCI 23, running between sites A and D needs additional bandwidth. When you reconfigure the network, you should notice several things when you baseline again. First will be the reduction in utilization as a percentage of CIR for DLCI 23. You should also notice less of a load on the virtual circuits providing an alternate path between these sites, as the primary path will not be as congested. Freeing up bandwidth on these alternate paths will in turn provide better performance for users that rely on them as their primary path to other sites. As you can see, the dynamics and interdependencies of “the cloud” are complex. A problem in one part of the cloud can have far reaching effects on other parts.

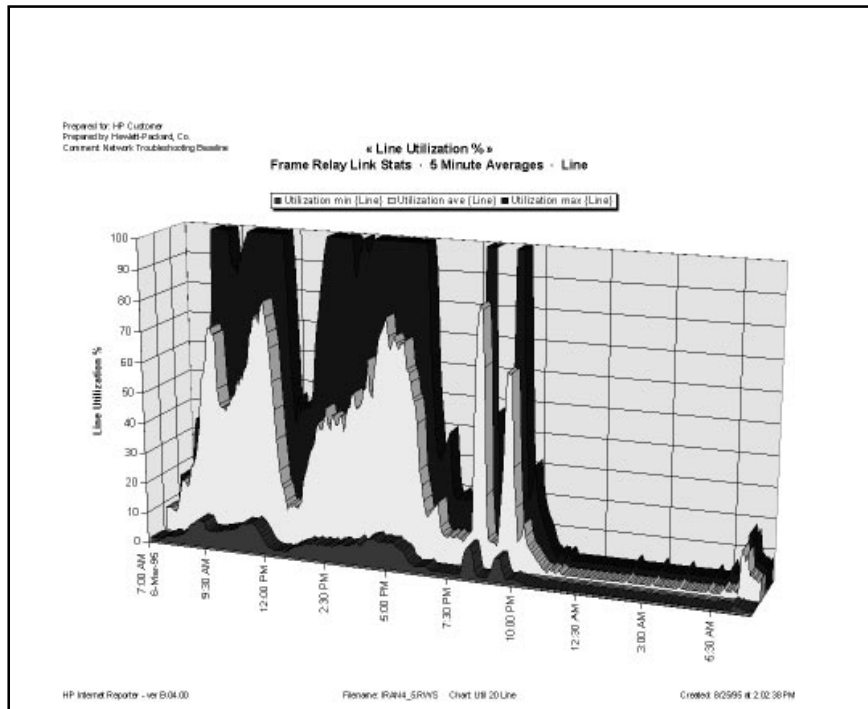


Figure 5.3 This chart shows the overall performance characteristics of the physical circuit at site A.




Figure 5.4 Using CIR values you enter, charts like this let you review individual virtual circuit bandwidth utilization.

Case Studies

Case 6: Verifying and documenting WAN circuit quality for trouble-free installations

You're in the process of relocating your data center and moving all of your WAN circuits into the new location. Although you've tried to maintain as much redundancy as possible during the move, you will only have three weeks of overlap before the old systems are taken off-line and you are completely dependent on the new data center. It is now three days before you bring up the first systems at the new facility, and your service provider tells you all the new circuits should be up and operational. Knowing that whatever can go wrong will go wrong, you need more assurance than someone else's word that the circuits and related communications equipment are operating reliably within specifications.

Internet Reporter solution

Combining the powerful bit error rate test (BERT) capabilities of the Internet Advisor with the reporting capabilities of Internet Reporter, you can easily verify and document the "quality" of WAN circuits and related communications equipment. BERT operates by injecting various test patterns into a network operating in loopback mode. Signals transmitted by the Internet Advisor are returned via the loopback and compared with the original patterns. Also measured when performing BERT are the various signal state changes and management codes transmitted by network devices, such as signal loss or yellow alarm. Each service provider publishes its own specifications for acceptable error rates based on the type

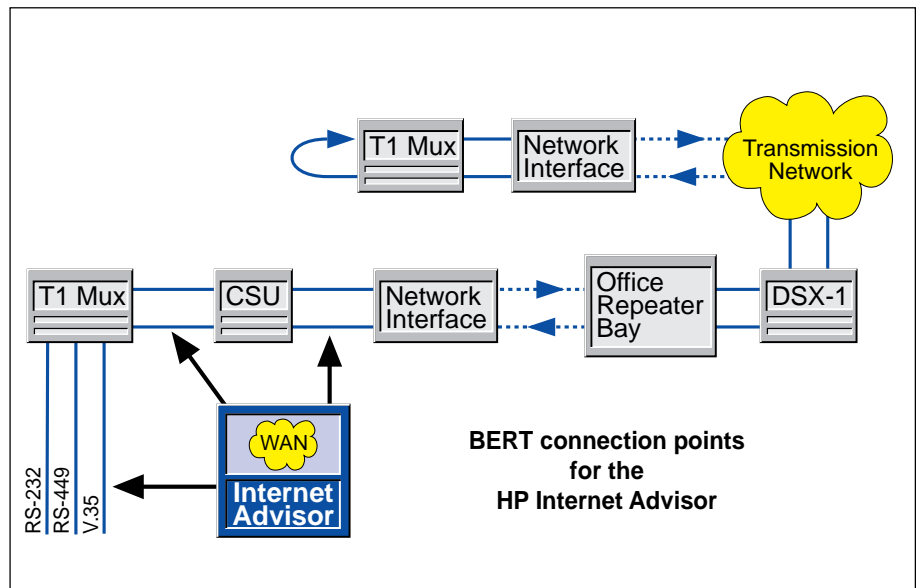


Figure 6.1 When performing BERT, you can connect your Internet Advisor directly to the network or through a V-Series interface.

of line, demark locations, and overall circuit distance. When reviewing Internet Reporter tables and charts, a quick visual scan will reveal conditions which are "out-of-spec" and need to be addressed by your service provider. In many cases, marginal quality circuits or circuits that are degrading go unnoticed until they degrade beyond the tolerance of your network equipment, and your network begins to fail. Performing BERT when the line is new and periodically checking for degradation go a long way toward providing reliable network services to your users.

Step one: plugging into the network

When connecting the Internet Advisor to the network, try to locate it as close to the demark as possible. (See figure 6.1.) If your

service provider allows it, plug directly into the T1 or E1 network. Be sure not to plug into the monitor jacks of the CSU or DSU, as this type of connection will not allow you to transmit BERT patterns on the network.

If it is not possible to connect directly to the network, connect to the V-series interface on the DSU. Although connecting directly to the network better isolates circuit-related problems, connecting through other communications equipment tests the additional cabling between devices and the devices themselves. For the first round of testing, you will connect directly to the network. Then, after you install your equipment, you will test again to ensure those devices and cables are also operating properly.

In addition to connecting an Internet Advisor at one end of the network, you will have to loopback the circuit to allow BERT to operate. For this first test of the circuit, there is probably no CSU installed on the far end, so you will have to ask your service provider to leave the far-end demark in loopback. Just remember to ask them to disable the loopback when you're done!

With all the connections made, you can run one of the BERT programs in the Internet Advisor Toolkit. Be sure to choose and configure a program that matches your circuit's configuration. Also be sure to enable disk logging before starting BERT so results may be imported into Internet Reporter. A suggested configuration for logging BERT statistics is to use a five to ten minute log

interval for a period of up to 24 hours. If you have more circuits to verify than time, you can spend as little as one hour on less critical circuits for basic verification.

Step two: charting and analyzing the results

As you complete each BERT session, you can copy the log file you created to disk and process it on a computer using Internet Reporter while you perform BERT on the next circuit. After processing the BERT log file with Internet Reporter, you'll have a concise table of BERT stats categorized in several efficiency, quality, and history sections (see figure 6.2.) Most of the statistics can be viewed graphically for easier analysis. Looking at the results of a test performed on this

sample circuit, you can see that this circuit needs to be repaired. (See figure 6.3.) According to the service provider's specifications, this 800 mile long circuit should have fewer than 50 errored seconds per day. Internet Reporter has documented several hundred errored seconds over the course of a day. Although your network equipment may be tolerant of these levels, it is indicative of potential problems at some time in the future. Now, before you install your equipment, you give your service provider a printed copy of your BERT documentation and ask for a rework the circuit to meet specifications.

After marginal or failing circuits are reworked by your service provider, retest each one until they are within specifications, then proceed to the next step.

Date		Basic Counts Per 10 Minutes				Basic Efficiency			G.821 Counts Per 10 Minutes				G.821 Efficiency %				T1 Errors Per 10 Minutes				T1 Events Per 10 Minutes										
Time	Bit count	Block count	Block errors	Error free seconds %	Bit error rate	Block error rate	Available time in seconds	Errored Seconds	Degraded Minutes	Severely errored seconds	Unavailable time in seconds	Available time	Errored Seconds	Degraded Minutes	Severely errored seconds	Unavailable time	Total T1 errors	ppv	Frame bit errors	Frame slip	ESF CRC errors	Signal loss	AS	Pulse density	Frame sync	Yellowalarm	BZS detect				
12:00 AM	45	6.97 E+8	1.38 E+6	100			45					75					115	73	90	1	11	1									
12:10 AM	600	9.38 E+8	1.83 E+6	100			600					100																			
12:20 AM	600	9.38 E+8	1.83 E+6	100			600					100																			
12:30 AM	600	9.38 E+8	1.83 E+6	100			600					100																			
12:40 AM	600	9.38 E+8	1.83 E+6	100			600					100																			
12:50 AM	600	9.38 E+8	1.83 E+6	100			600					100																			
1:00 AM	596	9.38 E+8	12 1.83 E+6	8	99	1.3 E-8	4.4 E-6	596	4	1	2	99	1	0	0	0															
1:10 AM	599	9.38 E+8	34 1.83 E+6	5	99	3.6 E-8	2.7 E-6	599	1	1	1	100	0	0	0	0															
1:20 AM	598	9.38 E+8	43 1.83 E+6	6	99	4.6 E-8	3.3 E-6	598	2	1	1	100	0	0	0	0															
1:30 AM	599	9.38 E+8	2 1.83 E+6	1	99	2.1 E-9	5.5 E-7	599	1	1	1	100	0	0	0	0															
1:40 AM	600	9.38 E+8	1.83 E+6	100			600					100																			
1:50 AM	600	9.38 E+8	1.83 E+6	100			600					100																			
2:00 AM	600	9.38 E+8	1.83 E+6	100			600					100																			
2:10 AM	589	9.38 E+8	43 1.83 E+6	15	99	4.6 E-8	8.2 E-6	589	11	1	6	3	98	2	0	1	1					1									
2:20 AM	598	9.38 E+8	32 1.83 E+6	6	99	3.4 E-8	3.3 E-6	598	2	1	1	100	0	0	0	0															
2:30 AM	600	9.38 E+8	1.83 E+6	100			600					100																			
2:40 AM	600	9.38 E+8	1.83 E+6	100			600					100																			
2:50 AM	600	9.38 E+8	1.83 E+6	100			600					100																			
3:00 AM	600	9.38 E+8	1.83 E+6	100			600					100																			
3:10 AM	600	9.38 E+8	1.83 E+6	100			600					100																			

Figure 6.2 This Internet Reporter table shows a wide range of BERT statistics that can help isolate potential network problems.

Step three: testing the rest of your communications equipment

Now that you have confidence in the quality of the lines, you can begin installing the rest of your communications equipment. Before allowing live data to pass through, however, perform BERT again on each circuit through the V-series, DSX-1, or CEPT/E1 test access ports of your equipment. Be sure to configure the far end loopbacks to include as much equipment inside the loop as possible. In the first series of tests you were checking circuit quality; now you are testing all the other equipment and cabling.

As you complete your tests, compare your results with the original circuit-only tests. If there are any variances, you will

have to troubleshoot your equipment and cabling to identify the problem component. As always, it can never hurt to retest the circuit to ensure that it has not developed problems since you first tested it.

Conclusion

Although nothing can guarantee that your network will operate properly all of the time, using the HP Internet Reporter to document BERT sessions can give you more peace of mind that your circuits and communications equipment are operating within specifications, and that any future problems are real failures, not the result of a bad installation.

HP Sales and Support Offices

For more information on Hewlett-Packard Test and Measurement products, applications or services please call your local Hewlett-Packard sales offices. A current listing is available via Web through AccessHP at <http://www.hp.com>. If you do not have access to the internet please contact one of the HP centers listed below and they will direct you to your nearest HP representative.

United States:

Hewlett-Packard Company
Test and Measurement Organization
5301 Stevens Creek Boulevard
Building 51L-SC
Santa Clara, CA 95052-8059
1 (800) 452-4844

Canada:

Hewlett-Packard Canada Ltd.
5150 Spectrum Way
Mississauga, Ontario
L4W 5G1
(905) 206-4725

Europe:

Hewlett-Packard
European Marketing Centre
P.O. Box 999
1180 AZ Amstelveen
The Netherlands

Japan:

Yokogawa-Hewlett-Packard Ltd.
Measurement Assistance Center
9-1, Takakura-Cho, Hachioji-Shi
Tokyo 192, Japan
(81) 426-48-3860

Latin America:

Hewlett-Packard
Latin American Region Headquarters
5200 Blue Lagoon Drive
9th Floor
Miami, Florida 33126
U.S.A.
(305) 267 4245/4220

Australia/New Zealand:

Hewlett-Packard Australia Ltd.
31-41 Joseph Street
Blackburn, Victoria 3130
Australia
131 347 ext. 2902

Asia Pacific:

Hewlett-Packard Asia Pacific Ltd.
17-21/F Shell Tower, Times Square,
1 Matheson Street, Causeway Bay,
Hong Kong
(852) 2599 7070

Data Subject to change

Printed in U.S.A. 9/95

© Copyright Hewlett-Packard
Company 1995



5964-2373E

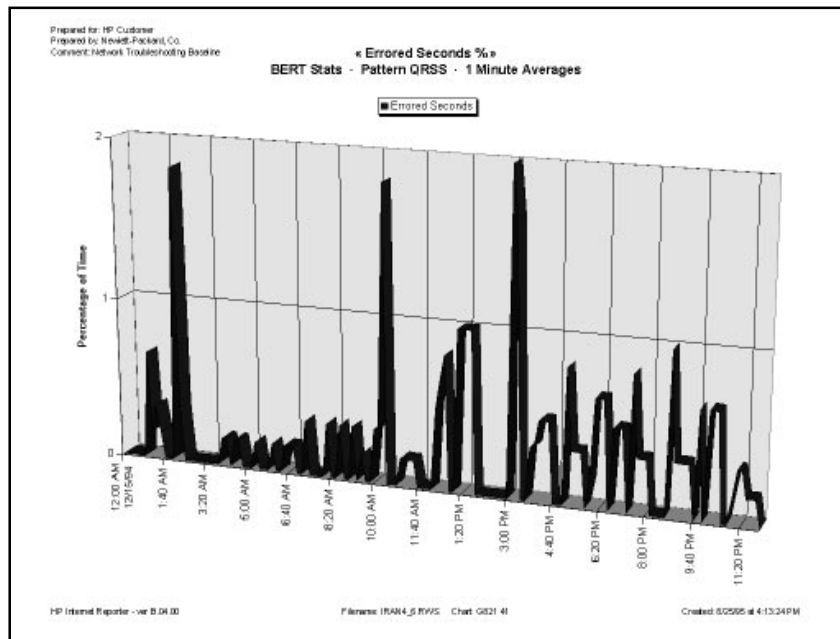


Figure 6.3 This BERT chart graphically details hundreds of errored seconds over a 24-hour period, indicating a line operating out-of-spec.