NAME

httperf – HTTP performance measurement tool

SYNOPSIS

httperf [−−**add−header** *R S*] [−−**burst−length** *R N*] [−−**client** *R I / N*] [−−**close−with−reset**] [−**d**|−−**debug** *R N*] [−−**failure−status** *R N*] [−**h**|−−**help**] [−−**hog**] [−−**http−version** *R S*] [−−**max−connections** *R N*] [−−**max−piped−calls** *R N*] [−−**method** *R S*] [−−**no−host−hdr**] [−−**num−calls** *R N*] [−−**num−conns** *R N*] [−−**period** [**d**|**u**|**e**]*T1*[,*T2*]] [−−**port** *R N*] [−−**print−reply** [**header**|**body**]] [−−**print−request** [**header**|**body**]] [−−**rate** *R X*] [−−**recv−buffer** *R N*] [−−**retry−on−failure**] [−−**send−buffer** *R N*] [−−**server** *R S*] [−−**server−name** *R S*] [−−**session−cookie**] [−−**ssl**] [−−**ssl−ciphers** *R L*] [−−**ssl−no−reuse**] [−−**think−timeout** *R X*] [−−**timeout** *R X*] [−−**uri** *R S*] [−**v**|−−**verbose**] [−**V**|−−**version**] [−−**wlog y**|**n,***F*] [−−**wsess** *R N , N , X*] [−−**wsesslog** *R N , X , F*] [−−**wset** *R N , X*]

DESCRIPTION

httperf is a tool to measure web server performance. It speaks the HTTP protocol both in its HTTP/1.0 and HTTP/1.1 flavors and offers a variety of workload generators. While running, it keeps track of a number of performance metrics that are summarized in the form of statistics that are printed at the end of a test run. The most basic operation of **httperf** is to generate a fixed number of HTTP GET requests and to measure how many replies (responses) came back from the server and at what rate the responses arrived.

IMPORTANT: To obtain correct results, it is necessary to run at most one **httperf** process per client machine. Also, there should be as few background processes as possible both on the client and server machines.

EXAMPLES

httperf −−hog −−server www
> This command causes **httperf** to create a connection to host www, send a request for the root document (http://www/), receive the reply, close the connection, and then print some performance statistics.

httperf −−hog −−server www −−num−conn 100 −−ra 10 −−timeout 5
> Like above, except that a total of 100 connections are created and that connections are created at a fixed rate of 10 per second. Note that option "−−rate" has been abbreviated to "−−ra".

httperf −−hog −−server=www −−wsess=10,5,2 −−rate 1 −−timeout 5
> Causes **httperf** to generate a total of 10 sessions at a rate of 1 session per second. Each session consists of 5 calls that are spaced out by 2 seconds.

httperf −−hog −−server=www −−wsess=10,5,2 −−rate=1 −−timeout=5 −−ssl
> Like above, except that **httperf** contacts server www via SSL at port 443 (the default port for SSL connections).

httperf −−hog −−server www −−wsess=10,5,2 −−rate=1 −−timeout=5 −−ssl −−ssl−ciphers=EXP−RC4−MD5:EXP−RC2−CBC−MD5 −−ssl−no−reuse −−http−version=1.0
> Like above, except that **httperf** will inform the server that it can only select from two cipher suites (EXP−RC4−MD5 or EXP−RC2−CBC−MD5); furthermore, **httperf** will use HTTP version 1.0 which requires a new TCP connection for each request. Also, SSL session ids are not reused, so the entire SSL connection establishment process (known as the SSL handshake) occurs for each connection.

OPTIONS

The operation of **httperf** can be controlled through a number of options. The tool supports both short (one−character) and long (arbitrary−length) option names. Short options are prefixed with a single leading dash (−), long options with a double−dash (−−). Multiple short options can be grouped together (e.g., "−**vV**" is equivalent to "−**v** −**V**") and long options can be abbreviated so long as they remain unique. Parameters to options can be specified either by following the long option name with an equal sign and the parameter value (e.g., −−**burst=10**) or by separating the option name and value with whitespace (e.g.,

−−burst 10).

−−add−header=*S*
> Specifies to include string *S* as an additional request header. It is necessary to specify the termi-
> nating carriage−return/line−feed sequence explicitly. This can be done by using the escape
> sequence "\n". This makes it possible to include multiple request headers. For example,
> "−−add−header "Referer: foo\nAuth: secret\n"" would add two request headers ("Referer" and
> "Auth") to each request. Other supported escape sequences are "\r" (carriage−return), "\a"
> (line−feed), "\\" (backslash), and "\N" where N is the code the character to be inserted (in octal).

−−burst−length=*N*
> Specifies the length of bursts. Each burst consists of *N* calls to the server. The exact meaning of
> this parameter depends on the workload generator. For regular request−oriented workloads, see
> the description of option **−−wsess**.

−−no−host−hdr
> Specifies that the "Host:" header should not be included when issuing an HTTP request.

−−num−calls.
> For session−oriented workloads, see the description of option **−−wsess**.

−−client=*I*/*N*
> Specifies that the machine **httperf** is running on is client *I* out of a total of *N* clients. *I* should be
> in the range from 0 to *R N −1*. Some of the workload generators (e.g., **−−wset)** use the client
> identity as a bias value to ensure that not all clients generate perfectly identical workloads. When
> performing a test that involves several client machines, it is generally a good idea to specify this
> option.

−−close−with−reset
> Requests that **httperf** closes TCP connections by sending a RESET instead of going through the
> normal TCP connection shutdown handshake. Turning on this option can have ill effects such as
> data corruption, stuck TCP control blocks, or wrong results. For this reason, the option should not
> be used unless absolutely necessary and even then it should not be used unless its implications are
> fully understood.

−d=*N*

−−debug=*N*
> Set debug level to *R N* . Larger values of *N* will result in more output.

−−failure−status=*N*
> Specifies that an HTTP response status code of *N* should be treated as a failure (i.e., treated as if
> the request had timed out, for example). For example, with "**−−failure−status=504**" responses
> with an HTTP status of "504 Gateway Time−out" would be considered failures. Caveat: this
> option is currently supported for session workloads only (see the **−−wsess** and **−−wsesslog**
> options).

−h

−−help Prints a summary of available options and their parameters.

−−hog This option requests to use up as many TCP ports as necessary. Without this option, **httperf** is
> typically limited to using ephemeral ports (in the range from 1024 to 5000). This limited port
> range can quickly become a bottleneck so it is generally a good idea to specify this option for seri-
> ous testing. Also, this option must be specified when measuring NT servers since it avoids a TCP
> incompatibility between NT and UNIX machines.

−−http−version=*S*
> Specifies the version string that should be included in the requests sent to the server. By default,
> version string "1.1" is used. This option can be set to "1.0" to force the generation of HTTP/1.0
> requests. Setting this option to any value other than "1.0" or "1.1" may result in undefined
> behavior.

−−max−connections=_N_

Specifies that at most _N_ connections are opened for each session. This option is meaningful in conjunction with options −−**wsess** and −−**wsesslog** only.

−−max−piped−calls=_N_

Specifies that at most _N_ pipelined calls are issued on each connection. This option is meaningful in conjunction with options −−**wsess** and −−**wsesslog** only.

−−method=_S_

Specifies the method that should be used when issuing an HTTP request. If this option is not specified, the GET method is used. The method _S_ can be an arbitrary string but is usually one of GET, HEAD, PUT, POST, etc.

−−num−calls=_N_

This option is meaningful for request−oriented workloads only. It specifies the total number of calls to issue on each connection before closing it. If _N_ is greater than 1, the server must support persistent connections. The default value for this option is 1. If −−**burst−length** is set to _R B_ , then the _N_ calls are issued in bursts of _B_ pipelined calls each. Thus, the total number of such bursts will be _N/B_ (per connection).

−−num−conns=_N_

This option is meaningful for request−oriented workloads only. It specifies the total number of connections to create. On each connection, calls are issued as specified by options −−**num−calls** and −−**burst−length**. A test stops as soon as the _N_ connections have either completed or failed. A connection is considered to have failed if any activity on the connection fails to make forward progress for more than the time specified by the timeout options −−**timeout** and −−**think−time-out**. The default value for this option is 1.

−−period=_[D]T1[,T2]_

Specifies the time interval between the creation of connections or sessions. Connections are created by default, sessions if option −−**wsess** or −−**wsesslog** has been specified. This connection/session ''interarrival time'' can alternatively be specified by the −−**rate** option, although more flexibility is available with −−**period.** The _D_ parameter specifies the interarrival time distribution. If omitted or set to ''**d**'', a deterministic (i.e., fixed) period is used as specified by parameter _R T1_ in units of seconds. If _D_ is set to ''**e**'', an exponential (i.e., Poisson) distribution is used with a mean interarrival time of _R T1_ . If _D_ is set to ''**u**'', a uniform distribution over the interval [_T1_,_T2_) is used for the interarrival time. Finally, if _D_ is set to ''**v**'', a number of rates can be specified as follows: −−**period=vT1,D1,T2,D2...Tn,Dn** Where n <= NUM_RATES in httperf.h and _Ti,Di_ represent the period time (i.e., 1/rate) and duration to maintain that rate (i.e., −−**period=v1,2,0.5,4** will generate 1 request/seconds for 2 seconds then 2 requests/seconds for 4 seconds). In all cases, a period of 0 results in connections or sessions being generated sequentially (a new connection/session is initiated as soon as the previous one completes). The default value for this option is 0. Note that specifying, for example, −−**rate=5** is equivalent to specifying −−**period=d0.2** or −−**period=0.2**. By specifying −−**period=u1,3**, the interarrival times will be randomly chosen from the interval between 1 and 3 seconds. The specific sequence of (pseudo−)random interarrival times are identical from one **httperf** run to another as long as the values for the −−**period** and −−**client** options are identical.

−−port=_N_

This option specifies the port number _N_ on which the web server is listening for HTTP requests. By default, **httperf** uses port number 80.

−−print−reply[=[header|body]]

Requests the printing of the reply headers, body, and summary. The output is directed to standard output. Reply header lines are prefixed by "RH", reply body lines are prefixed by "RB", and the reply−size summary is prefixed by "RS". The prefix is followed by a serial number that uniquely identifies the call that the reply line is for and a colon (":") character that marks the beginning of the actual reply line. To print only reply headers, pass argument **header** to this option. To print only the reply body, pass argument **body** to this option.

−−print−request[=[**header**|**body**]]

Requests the printing of the request headers, body (if one is present), and summary. The output is directed to standard output. Request header lines are prefixed by "SH", request body lines are prefixed by "SB", and the request summary is prefixed by "SS". The prefix is followed by the call's serial number and a colon (":") character that marks the beginning of the actual reply line. To print only request headers, pass argument **header** to this option. To print only the request body, pass argument **body** to this option.

−−rate=_X_

Specifies the fixed rate at which connections or sessions are created. Connections are created by default, sessions if option **−−wsess** or **−−wsesslog** has been specified. In both cases a rate of 0 results in connections or sessions being generated sequentially (a new session/connection is initiated as soon as the previous one completes). The default value for this option is 0.

−−recv−buffer=_N_

Specifies the maximum size of the socket receive buffers used to receive HTTP replies. By default, the limit is 16KB. A smaller value may help memory−constrained clients whereas a larger value may be necessary when communicating with a server over a high−bandwidth, high−latency connection.

−−retry−on−failure

This option is meaningful for session workloads only (see the **−−wsess** and **−−wsesslog** options). If specified, a call that results in a failure response (as defined by the **−−failure−status** option) is retried immediately instead of causing the session to fail.

−−send−buffer=_N_

Specifies the maximum size of the socket send buffers used to send HTTP requests. By default, the limit is 4KB. A smaller value may help memory−constrained clients whereas a larger value may be necessary when generating large requests to a server connected via a high−bandwidth, high−latency connection.

−−server=_S_

Specifies the IP hostname of the server. By default, the hostname "localhost" is used. This option should always be specified as it is generally not a good idea to run the client and the server on the same machine.

−−server−name=_S_

Specifies the (default) server name that appears in the "Host:" header of every request sent by **httperf**. Without this option, the host name (or IP address) specified by option **−−server** is used instead.

−−session−cookie

When this option is turned on, cookie managment is enabled on a per−session basis. What this means is that if a reply to a request that was generated by session _R X_ contains a cookie, then all future requests sent by session _X_ will include this cookie as well. At present, the cookie manager in **httperf** supports only one cookie per session. If a second cookie is received, the new cookie overwrites the existing one and a warning message is printed if ''−−debug 1'' is on.

−−ssl   Specifies that all communication between **httperf** and the server should utilize the Secure Sockets Layer (SSL) protocol. This option is available only if **httperf** was compiled with SSL support enabled.

−−ssl−ciphers=_L_

This option is only meaningful if SSL is in use (see **−−ssl** option). This option specifies the list _L_ of cipher suites that **httperf** may use in negotiating a secure connection with the server. If the list contains more than one cipher suite, the ciphers must be separated by a colon. If the server does not accept any of the listed cipher suites, the connection establishment will fail and **httperf** will exit immediately. If this option is not specified when the **−−ssl** option is present then **httperf** will use all of the SSLv3 cipher suites provided by the underlying SSL library.

−−ssl−no−reuse

This option is only meaningful if SSL and sessions are in use (see **−−ssl**, **−−wsess**, **−−wsesslog**). When an SSL connection is established the client receives a session identifier (session id) from the server. On subsequent SSL connections, the client normally reuses this session id in order to avoid the expense of repeating the (slow) SSL handshake to establish a new SSL session and obtain another session id (even if the client attempts to re−use a session id, the server may force the client to renegotiate a session). By default **httperf** reuses the session id across all connections in a session. If the **−−ssl−no−reuse** option is in effect, then **httperf** will not reuse the session id, and the entire SSL handshake will be performed for each new connection in a session.

−−think−timeout=X

Specifies the maximum time that the server may need to initiate sending the reply for a given request. Note that this timeout value is added to the normal timeout value (see option **−−timeout**). When accessing static web content, it is usually not necessary to specify this option. However, when performing tests with long−running CGI scripts, it may be necessary to use this option to allow for larger response−times. The default value for this option is zero seconds, meaning that the server has to be able to respond within the normal timeout value.

−−timeout=X

Specifies the amount of time X that **httperf** is willing to wait for a server reaction. The timeout is specified in seconds and can be a fractional number (e.g., **−−timeout 3.5**). This timeout value is used when establishing a TCP connection, when sending a request, when waiting for a reply, and when receiving a reply. If during any of those activities a request fails to make forward progress within the alloted time, **httperf** considers the request to have died, closes the associated connection or session and increases the **client−timo** error count. The actual timeout value used when waiting for a reply is the sum of this timeout and the think−timeout (see option **−−think−timeout**). By default, the timeout value is infinity.

−−uri=S

Specifies that URI S should be accessed on the server. For some of the workload generators (e.g., **−−wset**), this option specifies the prefix for the URIs being accessed.

−−use−timer−cache

This feature allows the user to specify whether they want to cache timestamps or not. Timestamps are not cached by default, but the user can enable caching if higher performance is more important than timing accuracy. For small response sizes, disabling timer caching reduced the performance of httperf by about 10%; for larger response sizes there was little or no effect.

−v

−−verbose

Puts **httperf** into verbose mode. In this mode, additional output such as the individual reply rate samples and connection lifetime histogram are printed.

−V

−−version

Prints the version of **httperf**.

−−wlog=B**,**F

This option can be used to generate a specific sequence of URI accesses. This is useful to replay the accesses recorded in a server log file, for example. Parameter F is the name of a file containing the ASCII NUL separated list of URIs that should be accessed. If parameter B is set to "**y**", **httperf** will wrap around to the beginning of the file when reaching the end of the list (so the list of URIs is accessed repeatedly). With B set to "**n**", the test will stop no later than when reaching the end of the URI list.

−−wsess=$N1$**,**$N2$**,**X

Requests the generation and measurement of sessions instead of individual requests. A session consists of a sequence of bursts which are spaced out by the user think−time. Each burst consists

of a fixed number *L* of calls to the server (*L* is specified by option −−**burst−length**). The calls in a burst are issued as follows: at first, a single call is issued. Once the reply to this first call has been fully received, all remaining calls in the burst are issued concurrently. The concurrent calls are issued either as pipelined calls on an existing persistent connection or as individual calls on separate connections. Whether a persistent connection is used depends on whether the server responds to the first call with a reply that includes a "Connection: close" header line. If such a line is present, separate connections are used.

The option specifies the following parameters: *N1* is the total number of sessions to generate, *N2* is the number of calls per session, and *X* is the user think−time (in seconds) that separates consecutive call bursts. For example, the options "−−**wsess=100,50,10** −−**burst−len 5**" would result in 100 sessions with a total of 50 calls each. Since each burst has a length of 5 calls, a total of 10 call bursts would be generated per session. The user think−time between call bursts would be 10 seconds. Note that user think−time *X* denotes the time between receiving the last reply of the previous call burst and the sending of the first request of the next burst.

A test involving sessions finishes as soon as the requested number *N1* of sessions have either failed or completed. A session is considered to have failed if any operation in a session takes longer than the timeouts specified by options −−**timeout** and −−**think−timeout**. In addition, a session also fails if the server returns a reply with a status code matching the one specified by option −−**failure−status**.

−−**wsesslog=***N***,***X***,***F*
> This specifies a session workload generator similar to −−**wsess** (please read that description first). With −−**wsesslog** though, many aspects of user sessions, including the number and sequence of URI's, request method, think−time and burst−length parameters, can be specified in an input file *F*. Two other parameters are retained from −−**wsess,** namely *N,* the number of sessions to initiate, and *X,* the burst−to−burst user think time (note that this becomes a default time since the input file *F* can also specify user think time on a per−burst basis. A small example input file can most−easily show the settable parameters:

> # Comment lines start with a "#" as the first
> # character. Lines with only whitespace delimit
> # sessions (multiple blank lines do not generate
> # "null" sessions). All other lines specify a
> # uri−sequence (1 uri per line). If the first
> # character of the line is whitespace (e.g. space
> # or tab), the uri is considered to be part of a
> # burst that is sent out after the previous
> # non−burst uri.

> # session 1 definition (this is a comment)
> /foo.html think=2.0
>   /pict1.gif
>   /pict2.gif
> /foo2.html method=POST contents='Post data'
>   /pict3.gif
>   /pict4.gif

> # session 2 definition
> /foo3.html method=POST contents="Multiline\ndata"
> /foo4.html method=HEAD

> The above description specifies 2 sessions. The first session will start with a request for /foo.html. When the /foo.html response comes back, a burst of 2 requests will follow (/pict1.gif and

/pict2.gif). When the last of those responses is received, a two second user think time is inserted before the next request of /foo2.html is issued. This request is sent as a POST. The posted data can be contained between single– or double–quotes. Newlines can appear within posted data as "\n" or as a "\<CR>". The /foo2.html response is followed by a burst request of /pict3.gif and /pict4.gif, which concludes this session. The second session is started some time after the first, as specified by the −−**rate** or −−**period** options.

The second session consists of 2 requests separated by the default user think time as specified by the *X* parameter of the −−**wsesslog** option. If the *N* parameter of −−**wsesslog** is greater than the number of sessions defined in input file *R F ,* then the defined sessions are used repeatedly until *N* sessions have been created (i.e., the defined sessions are used in a round–robin fashion).

One should avoid using −−**wsesslog** in conjunction with other **httperf** options that also control session behavior and workload URI's, namely −−**burst−length, −−wsess, −−wlog,** and −−**wset.**

−−**wset=***N,X*

This option can be used to walk through a list of URIs at a given rate. Parameter *N* specifies the number of distinct URIs that should be generated and *X* specifies the rate at which new URIs are accessed. A rate of **0.25** would mean that the same URI would be accessed four times in a row before moving on to the next URI. This type of access pattern is useful in generating a workload that induces a relatively predictable amount of traffic in the disk I/O subsystem of the server (assuming *N* and the accessed files are big enough to exceed the server's buffer cache). The URIs generated are of the form *R   prefix / path .html,* where *prefix* is the URI prefix specified by option −−**uri** and *path* is generated as follows: for the *R i −th* file in the working set, write down *i* in decimal, prefixing the number with as many zeroes as necessary to get a string that has as many digits as *R N −1.* Then insert a slash character between each digit. For example, the 103rd file in a working set consisting of 1024 files would result in a path of "**0/1/0/3**". Thus, if the URI–prefix is **/wset1024**, then the URI being accessed would be **/wset1024/0/1/0/3.html**. In other words, the files on the server need to be organized as a 10ary tree.

OUTPUT

This section describes the statistics output at the end of each test run. The basic information shown below is printed independent of the selected workload generator.

Total: connections 30000 requests 29997 replies 29997 test−duration 299.992 s

Connection rate: 100.0 conn/s (10.0 ms/conn, <=14 concurrent connections)
Connection time [ms]: min 1.4 avg 3.0 max 163.4 median 1.5 stddev 7.3
Connection time [ms]: connect 0.6
Connection length [replies/conn]: 1.000

Request rate: 100.0 req/s (10.0 ms/req)
Request size [B]: 75.0

Reply rate [replies/s]: min 98.8 avg 100.0 max 101.2 stddev 0.3 (60 samples)
Reply time [ms]: response 2.4 transfer 0.0
Reply size [B]: header 242.0 content 1010.0 footer 0.0 (total 1252.0)
Reply status: 1xx=0 2xx=29997 3xx=0 4xx=0 5xx=0

CPU time [s]: user 94.31 system 205.26 (user 31.4% system 68.4% total 99.9%)
Net I/O: 129.6 KB/s (1.1*10ˆ6 bps)

Errors: total 3 client−timo 0 socket−timo 0 connrefused 3 connreset 0
Errors: fd−unavail 0 addrunavail 0 ftab−full 0 other 0

There are six groups of statistics: overall results ("Total"), connection related results ("Connection"), results relating to the issuing of HTTP requests ("Request"), results relating to the replies received from the server ("Reply"), miscellaneous results relating to the CPU ("CPU") and network ("Net I/O") utilization and, last but not least, a summary of errors encountered ("Errors").

Total Section

This section summarizes how many TCP connections were initiated by **httperf**, how many requests it sent out, how many replies it received, and what the total test duration was. In the example output shown above, 30,000 connections were created, 29,997 requests were sent out and 29,997 replies were received. The duration of the test was almost exactly 5 minutes (300 seconds).

Connection Section

This section conveys information related to TCP connections generated by the tool. Specifically, the ''Connection rate'' line shows that new connections were initiated at a rate of 100.0 connections per second. This rate corresponds to a period of 10.0 milliseconds per connection. The last number in this line shows that at most 14 connections were open at any given time.

The first line labeled ''Connection time'' gives lifetime statistics for successful connections. The lifetime of a connection is the time between a TCP connection is initiated and the time the connection is closed. A connection is considered successful if it had at least one call that completed successfully. In the example output, the line indicates that the minimum (''min'') connection lifetime was 1.4 milliseconds, the average (''avg'') lifetime was 3.0 milliseconds, the maximum (''max'') was 163.4 milliseconds, the median (''median'') lifetime was 1.5 milliseconds, and that the standard deviation of the lifetimes was 7.3 milliseconds. The median lifetime is computed based on a histogram with one millisecond resolution and a maximum lifetime of 100 seconds. Thus, the median is accurate to within half a millisecond if at least half of the successful connections have a lifetime of no more than 100 seconds.

The next statistic in this section is the average time it took to establish a TCP connection. Only successful TCP connection establishments are counted. In the example, the second line labeled ''Connection time'' shows that, on average, it took 0.6 milliseconds to establish a connection.

The final line in this section is labeled ''Connection length.'' It gives the average number of replies received on each connection that received at least one reply (i.e., connections that failed before yielding the first reply are not counted). This number can be bigger than 1.0 due to persistent connections.

Request Section

The line labeled ''Request rate'' gives the rate at which HTTP requests were issued and the period that this rate corresponds to. In the example above, the request rate was 100.0 requests per second, which corresponds to 10.0 milliseconds per request. As long as no persistent connections are employed, the results in this section are very similar or identical to results in the connection section. However, when persistent connections are used, several calls can be performed on a single connection in which case the results would be different.

The line labeled ''Request size'' gives the average size of the HTTP requests in bytes. In the example above, the average request size was 75 bytes.

Reply Section

For simple measurements, this section is often the most interesting one as the line labeled ''Reply rate'' gives various statistics for the reply rate. In the example above, the minimum (''min'') reply rate was 98.8 replies per second, the average (''avg'') was 100 replies per second, and the maximum (''max'') rate was 101.2 replies per second. The standard deviation was 0.3 replies per second. The number enclosed in parentheses shows that 60 reply rate samples were acquired. At present, **httperf** collects a rate sample once every five seconds. To obtain a meaningful standard deviation, it is recommended to run tests long enough so at least thirty samples are obtained. This corresponds to a test duration of at least 150 seconds.

The line labeled ''Reply Time'' gives information on how long it took for the server to respond and how long it took to receive the reply. In the example, it took on average 2.4 milliseconds

between sending the first byte of the request and receiving the first byte of the reply. The time to "transfer", or read, the reply was too short to be measured, so it shows up as zero. The is typical when the entire reply fits into a single TCP segment.

The next line, labeled "Reply size" contains statistics on the average size of the replies−−−all numbers are in reported bytes. Specifically, the line lists the average length of reply headers, the content, and footers (HTTP/1.1 uses footers to realize the "chunked" transfer encoding). For convenience, the average total number of bytes in the replies is also given in parentheses. In the example, the average header length ("header") was 242 bytes, the average content length ("content") was 1010 bytes, and there were no footers ("footer" length is zero). The total reply length of 1252 bytes on average.

The final line in this section is a histogram of the major status codes received in the replies from the server. The major status code is the "hundreds"−digit of the full HTTP status code. In the example, all 29,997 replies had a major status code of 2. It's a good guess that all status codes were "200 OK" but the information in the histogram is not detailed enough to allow distinguishing status codes with the same major code.

Miscellaneous Section

This section starts with a summary of the CPU utilization on the client machine. In the example, the line labeled "CPU time" shows that 94.31 seconds were spent executing in user mode ("user"), 205.26 seconds were spent executing in system mode ("system") and that this corresponds to 31.4% user mode execution and 68.4% system execution. The total utilization was 99.9%, which is expected given that **httperf** is a CPU hog. A total CPU utilization of significantly less than 100% is a sign that there were competing processes that interfered with the test.

The line labeled "Net I/O" gives the average network throughput in kilobytes per second (where a kilobyte is 1024 bytes) and in megabits per second (where a megabit is 10ˆ6 bits). In the example, an average network usage of about 129.6 kilobytes per second was sustained. The number in parentheses shows that this corresponds to about 1.1 megabits per second. This network bandwidth is computed based on the number of bytes sent and received on the TCP connections. In other words, it does not account for the network headers or TCP retransmissions that may have occurred.

Errors Section

The last section contains statistics on the errors that were encountered during a test. In the example, the two lines labeled "Errors" show that there were a total of three errors and that all three errors were due to the server refusing to accept a connection ("connrefused"). A description of each error counter follows:

client−timo: The number of times a session, connection, or call failed due to a client timeout (as specified by the **−−timeout** and **−−think−timeout**) options.

socket−timo: The number of times a TCP connection failed with a socket−level timeout (ETIMEDOUT).

connrefused: The number of times a TCP connection attempt failed with a "connection refused by server" error (ECONNREFUSED).

connreset: The number of times a TCP connection failed due to a RESET from the server. Typically, a RESET is received when the client attempts to send data to the server at a time the server has already closed its end of the connection. NT servers also send RESETs when attempting to establish a new connection when the listen queue is full.

fd−unavail: The number of times the **httperf** process was out of file descriptors. Whenever this count is non−zero, the test results are meaningless because the client was overloaded (see section

"CHOOSING TIMEOUT VALUES").

addrunavail: The number of times the client was out of TCP port numbers (EADDRNOTAVAIL). This error should never occur. If it does, the results should be discarded.

ftab−full: The number of times the system's file descriptor table is full. Again, this error should never occur. If it does, the results should be discarded.

other: The number of times some other type of error occurred. Whenever this counter is non−zero, it is necessary to track down the real cause of the error. To assist in doing this, **httperf** prints the error code (errno) of the first unknown errors that occurs during a test run.

When −−**wsess** or −−**wsesslog** is specified, **httperf** generates and measures sessions instead of individual calls and additional statistics are printed at the end of a test. An example output is shown below.

Session rate [sess/s]: min 0.00 avg 0.59 max 2.40 stddev 0.37 (240/450)
Session: avg 6.45 connections/session
Session lifetime [s]: 123.9
Session failtime [s]: 58.5
Session length histogram: 4 7 4 ... 3 3 240

The line labeled ''Session rate'' shows the minium, average, and maximum rate at which sessions completed (based on a 5 second sampling interval). It also shows the standard deviation of the session completion rate. The numbers in parentheses show how many sessions succeeded and how many sessions were initiated. In the example above, the minimum, average, and maximum session completion rates were 0.00, 0.59, and 2.40 sessions per second, respectively. The standard deviation was 0.37 sessions per second and 240 out of 450 sessions completed successfully (210 failed due to errors such as timeouts).

The next line, labeled ''Session:'' shows the average length of a session measured in connections. In the example above, an average of 6.45 connections were required to complete a session.

The line labeled ''Session lifetime'' gives the average time it took to complete a successful session. In the example above, it took an average of 123.9 seconds.

The line labeled ''Session failtime'' gives the average time it took before an unsuccessful session failed. In the example above, it took on average 58.5 seconds for a session to fail.

Finally, the line labeled ''Session length histogram'' gives a histogram of the number of replies received by each session. In the example above, 4 sessions ended after receiving no reply at all, 7 ended after receiving one reply, and so on (the ellipsis indicates additional histogram counts that were omitted from this manual for space reasons). Note that this histogram does not distinguish between successful and failed sessions.

CHOOSING TIMEOUT VALUES

Since the machine that **httperf** runs on has only a finite set of resource available, it can not sustain arbitrarily high HTTP loads. For example, one limiting factor is that there are only roughly 60,000 TCP port numbers that can be in use at any given time. Since on most UNIX systems it takes one minute for a TCP connection to be fully closed (leave the TIME_WAIT state), the maximum rate a client can sustain is at most 1,000 requests per second.

The actual sustainable rate is often much lower than that because before running out of TCP ports, the machine is likely to run out of file descriptors (one file descriptor is used up for each open TCP connection). By default, HP−UX 10.20 allows 1,024 open file descriptors per process. This means that without extra precautions, **httperf** could potentially very quickly use up all available file descriptors, at which point it could not induce any additional load on the server. To avoid this problem, **httperf** provides option −−**timeout** to set a timeout for all communication with the server. If the server does not respond before the timeout expires, the client considers the corresponding session, connection, or call to be ''dead,'' closes the

associated TCP connection, and increases the "client−timo" error count. The only exception to this rule is that after sending an entire request to the server, **httperf** allows the server to take some additional time before it starts sending the reply. This is to accommodate HTTP requests that take a long time to complete on the server. This additional time is called the "server think time" and can be specified by option −−**think−timeout**. By default, this additional think time is zero seconds, so the server would always have to respond within the time alloted by option −−**timeout**.

Timeouts allow **httperf** to sustain high offered loads even when the server is overloaded. For example, with a timeout of 2 seconds and assuming that 1,000 file−descriptors are available, the offered load could be up to 500 requests per second (in practice, the sustainable load is often somewhat smaller than the theoretical value). On the downside, timeouts artificially truncate the connection lifetime distribution. Thus, it is recommended to pick a timeout value that is as large as possible yet small enough to allow sustaining the desired offered rate. A timeout as short as one second may be acceptable, but larger timeouts (5−10 seconds) are preferable.

It is important to keep in mind that timeouts do not guarantee that a client can sustain a particular offered load−−−there are many other potential resource bottlenecks. For example, in some cases the client machine may simply run out of CPU time. To ensure that a given test really measured the server's capabilities and not the client's, it is a good idea to vary the number of machines participating in a test. If observed performance remains the same as the number of client machines is varied, the test results are likely to be valid.

AUTHOR

httperf was developed by David Mosberger and was heavily influenced by an earlier tool written by Tai Jin. Stephane Eranian contributed the log−file based URI generator. Dick Carter contributed the −−**wsess-log** workload generator, the support behind the −−**period** option, and bug fixes. All four authors are with Hewlett−Packard Research Laboratories.

BUGS

Probably many. Always be sure to double−check results and don't fall prey to measuring client−performance instead of server performance!

The user−interface definitely could be improved. A simple workload description language might be more suitable than the dozens of little command−line options the tool has right now.