

Performance Metrics & Analysis

Network Management Workshop

APRICOT 2009

18 February 2009

Hervey Allen / Phil Regnauld

Original Materials in Spanish by Carlos Vicente, University of Oregon Network Services



Contents

- Planning performance management
- Metrics
 - Network
 - Systems
 - Services
- Measurement examples

Planning

- What's the intention?
 - *Baselining, Troubleshooting, Planning growth*
 - Defend yourself from accusations - "it's the network!"
- Who is the information for?
 - Administration, NOC, customers
 - How to structure and present the information
- Reach: Can I measure everything?
 - Impact on devices (measurements and measuring)
 - Balance between amount of information and time to get it

Metrics

- Network performance metrics
 - Channel capacity, nominal & effective
 - Channel utilization
 - Delay and *jitter*
 - Packet loss and errors
- System performance metrics
 - Availability
 - Memory, CPU Utilization, *load*, *I/O wait*, etc.
- Service performance metrics

Common network performance measurements

- Relative to traffic:
 - Bits per second
 - Packets per second
 - *Unicast vs. non-unicast* packets
 - Errors
 - Dropped packets
 - Flows per second
 - Round trip time (RTT)
 - Jitter (variation between packet RTT)

Nominal channel capacity

- The maximum number of bits that can be transmitted for a unit of time (eg: bits per second)
- Depends on:
 - Bandwidth of the physical medium
 - Cable
 - Electromagnetic waves
 - Processing capacity for each transmission element
 - Efficiency of algorithms in use to access medium
 - Channel encoding and compression

Effective channel capacity

- Always a fraction of the nominal channel capacity
- Dependent on:
 - Additional overhead of protocols in each layer
 - Device limitations on both ends
 - Flow control algorithm efficiency, etc.
 - For example: TCP


Channel utilization

- What fraction of the nominal channel capacity is actually in use
- Important!
 - Future planning
 - What utilization growth rate am I seeing?
 - For when should I plan on buying additional capacity?
 - Where should I invest for my updates?
 - Problem resolution
 - Where are my bottlenecks, etc.


95th Percentile

- The smallest value that is larger than 95% of the values in a given sample
- This means that 95% of the time the channel utilization is equal to or *less* than this value
 - Or rather, the peaks are discarded from consideration
- Why is this important in networks?
 - Gives you an idea of the standard, sustained channel utilization.
 - ISPs use this measure to bill customers with “larger” connections.

95th Percentile




Bits per second vs Packets p.s.



End-to-end delay

- The time required to transmit a packet along its entire path
 - *Created by an application, handed over to the OS, passed to a network card (NIC), encoded, transmitted over a physical medium (copper, fibre, air), received by an intermediate device (switch, router), analyzed, retransmitted over another medium, etc.*
 - The most common measurement uses *ping* for total round-trip-time (RTT).

Historical measurement of delay



Types of Delay

- Causes of end-to-end delay
 - Processor delays
 - Buffer delays
 - Transmission delays
 - Propagation delays

Processing delay

- Required time to analyze a packet header and decide where to send the packet (eg. a routing decision)
 - Inside a router this depends on the number of entries in the routing table, the implementation of data structures, hardware in use, etc.
- This can include error verification / checksumming (i.e. IPv4, IPv6 header checksum)

Queuing Delay

- The time a packet is enqueued until it is transmitted
- The number of packets waiting in the queue will depend on traffic intensity and of the type of traffic
- Router queue algorithms try to adapt delays to specific preferences, or impose equal delay on all traffic.

Transmission Delay

- The time required to push all the bits in a packet on the transmission medium in use
- For N =Number of bits, S =Size of packet, d =delay

$$d = S/N$$

- For example, to transmit 1024 bits using Fast Ethernet (100Mbps)

$$d = 1024/1 \times 10^8 = 10.24 \text{ micro seconds}$$

Propagation Delay

- Once a bit is 'pushed' on to the transmission medium, the time required for the bit to propagate to the end of its physical trajectory
- The velocity of propagation of the circuit depends mainly on the actual distance of the physical circuit
 - In the majority of cases this is close to the speed of light.
- For d = distance, s = propagation velocity

$$PD = d/s$$



Transmission vs. Propagation

- Can be confusing at first
- Consider this example:
 - Two 100 Mbps circuits
 - 1 km of optic fiber
 - Via satellite with a distance of 30 km between the base and the satellite
 - For two packets of the same size which will have the larger transmission delay? Propagation delay?

Packet Loss

- Occur due to the fact that buffers are not infinite in size
 - When a packet arrives to a buffer that is full the packet is discarded.
 - Packet loss, if it must be corrected, is resolved at higher levels in the network stack (transport or application layers)
 - Loss correction using retransmission of packets can cause yet more congestion if some type of (flow) control is not used (to inform the source that it's pointless to keep sending more packets at the present time)

Jitter



Flow Control and Congestion

- Limits the transmission amount (rate) because the receiver cannot process packets at the same rate that packets are arriving.
- Limit the amount sent (transmission rate) because of loss or delays in the circuit.

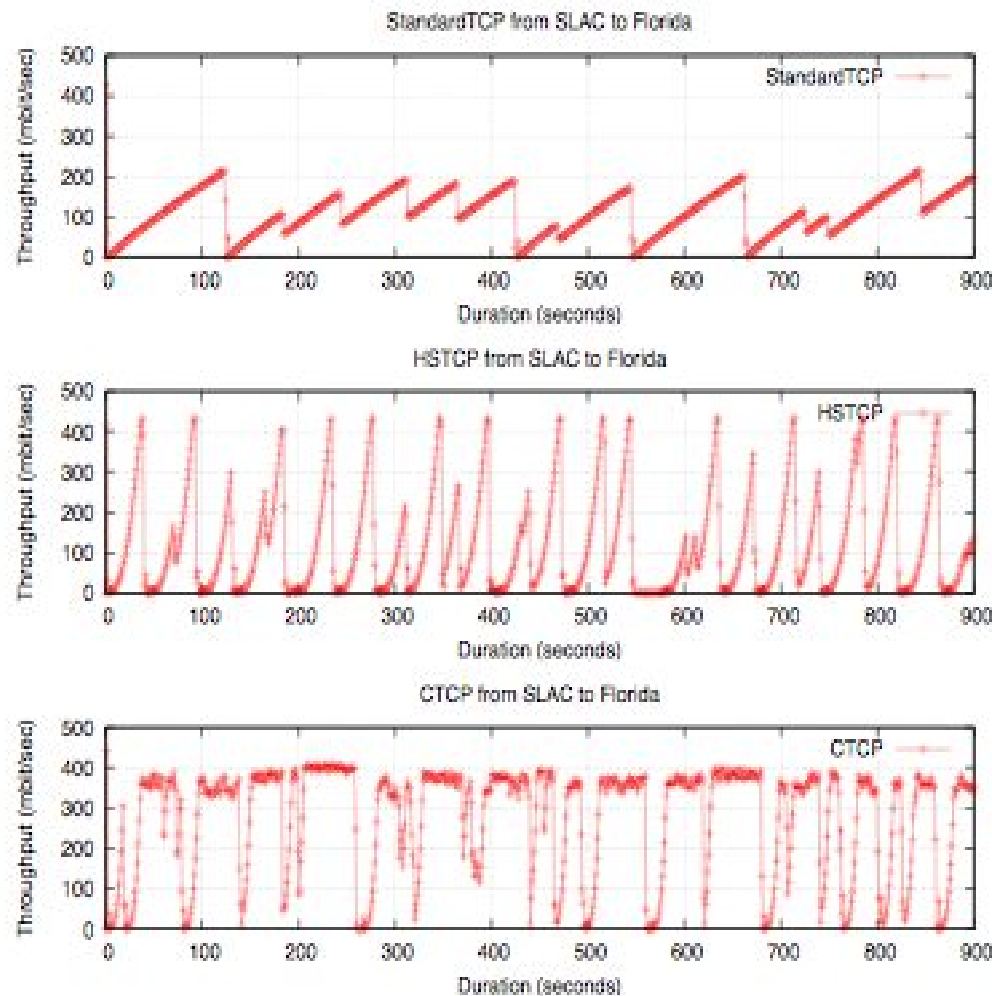
Controls in TCP

- IP (Internet Protocol) implements service that not connection oriented.
 - There is no mechanism in IP to deal with packet loss.
- TCP (Transmission *Control* Protocol) implements flow and congestion control.
 - Only on the ends as the intermediate nodes at the network level do not talk TCP

Congestion vs. Flow in TCP

- Flow: controlled by window size (RcvWindow), which is sent by the receiving end.
- Congestion: controlled by the value of the congestion window (Congwin)
 - Maintained independently by the sender
 - This varies based on the detection of packets lost
 - Timeout or receiving three ACKs repeated
 - Behaviors:
 - Additive Increments / Multiplicative Decrements (AIMD)
 - Slow Start
 - React to *timeout* events

Different TCP Congestion Control Algorithms



Systems Measurements

- Availability
- Unix/Linux Systems:
 - CPU usage
 - Kernel, System, User, IOwait
 - Memory usage
 - Real and Virtual
 - Load


Availability

State	Type / Reason	Time	% Total Time	% Known Time
UP	Unscheduled	363d 15h 35m 6s	99.995%	99.995%
	Scheduled	1d 8h 0m 45s	0.365%	0.365%
	Total	364d 23h 35m 51s	99.995%	99.995%
DOWN	Unscheduled	0d 0h 0m 0s	0.000%	0.000%
	Scheduled	0d 0h 0m 0s	0.000%	0.000%
	Total	0d 0h 0m 0s	0.000%	0.000%
UNREACHABLE	Unscheduled	0d 0h 23m 36s	0.005%	0.005%
	Scheduled	0d 0h 0m 33s	0.000%	0.000%
	Total	0d 0h 24m 9s	0.005%	0.005%
Undetermined	Nagios Not Running	0d 0h 0m 0s	0.000%	
	Insufficient Data	0d 0h 0m 0s	0.000%	
	Total	0d 0h 0m 0s	0.000%	
All	Total	365d 0h 0m 0s	100.000%	100.000%


State Breakdowns For Host Services:

Service	% Time OK	% Time Warning	% Time Unknown	% Time Critical	% Time Undetermined
HTTP	99.999% (99.999%)	0.000% (0.000%)	0.000% (0.000%)	0.001% (0.001%)	0.000%
HTTPS	99.974% (99.974%)	0.000% (0.000%)	0.000% (0.000%)	0.026% (0.026%)	0.000%
PING	0.000% (0.000%)	0.000% (0.000%)	0.000% (0.000%)	0.000% (0.000%)	100.000%
POP3	99.926% (99.926%)	0.000% (0.000%)	0.000% (0.000%)	0.074% (0.074%)	0.000%
SMTP	99.998% (99.998%)	0.000% (0.000%)	0.000% (0.000%)	0.002% (0.002%)	0.000%
TRAP	25.865% (25.865%)	0.000% (0.000%)	0.000% (0.000%)	74.135% (74.135%)	0.000%
Average	70.960% (70.960%)	0.000% (0.000%)	0.000% (0.000%)	12.373% (12.373%)	16.667%


CPU Usage



Memory





System load (I/O / CPU wait states)




Measuring services

- The key is to choose the most important measurements for each service
 - Ask yourself:
 - How is service degradation perceived
 - Wait time / Delay
 - Availability?
 - How can I justify maintaining the service?
 - Who is using it?
 - How often?
 - Economic value? Other value?


Web server usage



Response Time (Web server)






Response Time (DNS Server)



DNS Measurements

Result	Description
Success	Number of queries that resulted in a success (not a referral)
Referral	Number of queries that resulted in referrals
NXRRSET	Number of queries that resulted in a non-existent requested Resource Record Set
NXDOMAIN	Number of queries where the queried name does not exist
Recursion	Number of queries that required the sending of additional queries to the server
Failure	Number of queries that resulted in errors other than NXDOMAIN (serv fail, ...)
Total	Number of queries by unit of time


DNS Measurements



Mail Server Statistics

- Counters by *mailer* (local, SMTP, etc.)
 - Number of received/sent messages
 - Number of received/sent bytes
 - Number of rejected messages
 - Number of dropped messages
- Very important: number of queued messages
- Delivery rate
- Direction (inbound, outbound, inside, outside)




Sendmail Statistics




Web Proxy Measurements

- Number of requests per seconds
- Requests served locally vs. those requested externally
 - Web destination diversity
 - Efficiency of our web proxy
- Number of elements stored on disk vs. in memory

Squid Statistics



DHCP Statistics



Questions ?