SHARKFEST '13

Wireshark Developer and User Conference

Deep Dive Packet Analysis

Hansang Bae, Director
Riverbed Performance Management, Architect

http://www.youtube.com/hansangb has the Camtasia recorded sessions.
https://www.box.com/Sharkfest2013 has the trace files used in this session
REFER to the TAKE-AWAY sheet at the end for some key notes/findings!

Hidden Dangers of DC Migration

- Before any data center migration is performed, you must identify the application ecosystem.
  - If you fail to do this, prepare to update your resume/CV!
  - Application diagrams have magical lines. But then again, network diagrams are deaf, mute, and blind to applications.
  - Application flow diagram is a good start. However, you must use the “trust but verify” approach. If not, see bullet #1 above.
  - Netflow is your only chance of getting this right.

Hidden Dangers of DC Migration

Dangers Nagle/Delayed Ack with a twist!

- TCP was never meant for real time applications.
  - If using small packets, that is.
  - Even the handshake takes three packets!
  - Do you *really* know the intricacies of Nagle and Delayed Ack?
  - Turning off Delayed Ack or Nagle is a cop out. But it's easier than rewriting the entire application.
  - This application uses a 3 byte "start of field" marker, followed by additional data.
  - Datacenter is 11ms apart.

File Transfer is slow in one direction

- Remember what's important in file transfer throughput issues?
  - You have to rule out Window size issues.
  - Retransmissions (not fast retransmissions) that cause slow start behavior – draining the pipe.
  - Is buffer tearing going on? More common than you think.

THANK YOU!!!

- I think I'd be remiss if I didn't thank Gerald, Janice, all the core developers and fellow presenters.
- Special thanks to Rich Siefert and Charles Kaplan for the outstanding key note speeches.
- It boggles my mind how much I learn every time I attend Sharkfest.
- If you consider yourself a protocol analyst, performance engineer, level 3 operations, or network troubleshooter, than you really must attend Sharkfest.
- Finally, any and all feedback (good or bad, but especially bad) is welcome. If you don't tell me what you didn't like, I can't fix it!!! See you at Sharkfest 2014!

Takeaway Sheet

- File Transfer issues
  - Use PSH bits to identify if application is trying to dictate the transfer. We call this buffer tearing. Remember, lot of PSH bits in the trace means “danger, danger, Will Robinson” So be very careful about migrating applications with lot of PSH bits. In fact, DON’T MIGRATE applications that fall into this bucket.
  - Keep in mind that you can only transfer one full TCP window size per round-trip. Learn it, Know it, Live it!
  - If you don’t know EXACTLY how TCP works, you’ll never know that “delay” seen between packets 10 and 11 (NagleRemovedFromOneSide.pcap) is impossible. That TCP is *not* at fault for this delay.
  - Why is that delay not possible? Since Nagle has been delayed, the sender is free to send two small – non full MSS – packets without having to wait.
  - Refer to my “Hidden Dangers of Nagle Delayed Ack” on my youtube channel (http://www.youtube.com/hansangb) for more about Nagle.