

Effective Diagnostic Strategies for Wide Area Networks
aka
Network Path and Application Diagnosis (NPAD)

Peter O'Neil

NCAB

September 15, 2004

NSF STI Award

- 3 year PSC & NCAR project
 - Matt Mathis (PSC) & Peter O’Neil (NCAR) Co-PI’s
 - Senior Personnel
 - John Heffner & Raghu Reddy from PSC
 - Pete Siemsen & David Mitchell from NCAR
- Builds on Web100 and Net100 projects
- Continues focus on e2e performance issues
 - where previous efforts dealt with end systems, we now turn our attention to application flows across the path

Observed Behavior



- Symptoms appear to scale with increasing RTT path delay
 - Servers and Local Clients throughput good, but for remote client, throughput performance is poor
 - Reflective of most types of flaws
 - Impacts are multiplicative as backbone path magnifies symptoms of an existing flaw

Example Impacts

- Chat application (e.g., 50 RTT per user request)
 - On 1ms LAN, 50ms total time
 - On 100ms WAN, 5s total time
- Fixed TCP buffer space (e.g., 32kBytes)
 - On a 1ms LAN, 200Mb/s throughput limit
 - On a 100ms WAN, 2Mb/s throughput limit
- Packet Loss (e.g., 1% with 9kB MTU packets)
 - On a 1ms LAN, 500 Mb/s throughput limit
 - On a 100ms WAN, 5Mb/s throughput limit


Symptoms Present Diagnostic Quandaries

- False reassurance on short path flows
 - Obscures actual local bugs/problems
 - Incorrectly points elsewhere
- Stymies e2e diagnosis and thus improvements
- Promotes diagnosis as “one off workarounds”
 - Unscalable time sync

NPAD Approach

- Web100 based diagnostic server
 - Simple TCP test to a test target
 - Use MIB and model to rescale results
- TCP discard server for test target
 - Believe (hope) would be trivial to widely deploy
 - C or Java
- Provide good estimates of results for long paths

e2e Diagnosis



- Approximately 1 diagnostic server per campus/backbone/GigaPoP
- Test targets at lots of PoPs/hubs or LANs (or WS/clusters)
- Allows for isolating a flaw (slowdown) across each element of a long path

End User Tool

- Web/Java client
 - With built in test target
 - Invokes test on Diagnostic Server back to self

Combine with Other Approaches

- Bench test applications and end-systems (stacks)
- Use long ideal (virtual) paths
 - Dummynet style of emulated delay
 - Tunnel or VPN style of “scenic” routing

Project Goals

- Develop and learn from tool extensions which compensate for results that scale with RTT
- Study effects of various delay times
- Test the effects of these diagnostic tools with network users and operators using actual high performance applications

NPAD Summary Rationale

- Single point failures are (relatively) easy to find and fix
- Remaining failures are interactive & complex
 - RTT, packet loss, and MTU size
 - RTT and application design
 - Packet rate limit and MTU
 - Queue size buffers in routers and burstyness due to
 - ACK compression and cross traffic or
 - Application design