



August 10th 2010, OSG Site Admin Workshop - Network Performance
Jason Zurawski, Internet2

Welcome & Performance Primer

Who am I, Who are you?

Agenda

- Welcome and Thanks
 - <http://www.internet2.edu/workshops/npw/>
- Tutorial Agenda:
 - Network Performance Primer - Why Should We Care? **(15 Mins)**
 - Getting the Tools **(10 Mins)**
 - Use of the BWCTL Server and Client **(30 Mins)**
 - Use of the OWAMP Server and Client **(30 Mins)**
 - Use of the NDT Server and Client **(30 Mins)**
 - BREAK **(15 mins)**
 - Diagnostics vs Regular Monitoring **(30 Mins)**
 - Network Performance Exercises **(1 hr 30 Mins)**

Your Goals?

- What are your goals for this workshop?
 - Experiencing performance problems?
 - **Responsible for the campus/lab network?**
 - Learning about state of the art, e.g. ‘What is perfSONAR’?
 - Developing or researching performance tools?
- Is there a Magic Bullet?
 - No, but we can give you access to strategies and tools that will help
 - Patience and diligence will get you to most goals
- This workshop is as much a learning experience for me as it is for you
 - What problem/problems need to be solved
 - What will make networking a less painful experience
 - How can we improve our goods/services

Problem: “The Network Is Broken”

- How can your users effectively report problems?
- How can you users and the local administrators effectively solve multi-domain problems?
- Components:
 - Tools to use
 - Questions to ask
 - Methodology to follow
 - How to ask for (and receive) help

Why Worry About Network Performance?

- Most network design lends itself to the introduction of flaws:
 - Heterogeneous equipment
 - Cost factors heavily into design – e.g. *Get what you pay for*
 - Design heavily favors **protection** and **availability** over performance
- Communication protocols are not advancing as fast as networks
 - *TCP/IP* is the king of the protocol stack
 - Guarantees reliable transfers
 - Adjusts to failures in the network
 - Adjusts speed to be *fair* for all
- User Expectations
 - **Big Science** is prevalent globally
 - “The Network is Slow/Broken” – is this the response to almost any problem? Hardware? Software?
 - Empower users to be more informed/more helpful

Motivation – A Typical Scenario

- User and resource are geographically separated
- Both have access to high speed communication network
 - LAN infrastructure - 1Gbps Ethernet
 - WAN infrastructure – 10Gbps Optical Backbone



Motivation – A Typical Scenario

- User wants to access a file at the resource (e.g. ~600MB)
- Plans to use COTS tools (e.g. “scp”, but could easily be something scientific like “GridFTP” or simple like a web browser)
- What are the expectations?
 - 1Gbps network (e.g. *bottleneck* speed on the LAN)
 - $600\text{MB} * 8 = 4,800 \text{ Mb}$ file
 - User expects *line rate*, e.g. $4,800 \text{ Mb} / 1000 \text{ Mbps} = 4.8 \text{ Seconds}$
 - Audience Poll: Is this expectation too high?
- What are the realities?
 - Congestion and other network performance factors
 - Host performance
 - Protocol Performance
 - Application performance

Motivation – A Typical Scenario

- Real Example (New York USA to Los Angeles USA):

```
[zurawski@nms-rthr2 ~]$ scp zurawski@bwctl1.losa.net.internet2.edu:pS-Performance_Toolkit-3.1.1.iso .  
pS-Performance_Toolkit-3.1.1.iso          2%  17MB  1.0MB/s  10:05 ETA_
```

• Example:

- 1MB/s (8Mb/s) ??? 10 Minutes to transfer???
- Seems unreasonable given the investment in technology
 - Backbone network
 - High speed LAN
 - Capable hosts
- Performance realities as network speed decreases:
 - 100 Mbps Speed – 48 Seconds
 - 10 Mbps Speed – 8 Minutes
 - 1 Mbps Speed – 80 Minutes
- How could this happen? More importantly, why are there not more complaints?
- Audience Poll: Would you complain? If so, to whom?
- Brainstorming the above – where should we look to fix this?

Motivation – A Typical Scenario

- Expectation does not even come close to experience, time to debug.
Where to start though?
 - Application
 - Have other users reported problems? Is this the most up to date version?
 - Protocol
 - Protocols typically can be tuned on an individual basis, consult your operating system.
 - Host
 - Are the hardware components (network card, system internals) and software (drivers, operating system) functioning as they should be?
 - LAN Networks
 - Consult with the local administrators on status and potential choke points
 - Backbone Network
 - Consult the administrators at remote locations on status and potential choke points (Caveat – do you [should you] know who they are?)

Motivation – A Typical Scenario

- Following through on the previous, what normally happens ...
 - Application
 - This step is normally skipped, the application designer will *blame the network*
 - Protocol
 - These settings may not be explored. Shouldn't this be automatic (e.g. autotuning)?
 - Host
 - Checking and diagnostic steps normally stop after establishing connectivity. E.g. “can I ping the other side”
 - LAN Networks
 - Will assure “internal” performance, but LAN administrators will ignore most user complaints and shift blame to upstream sources. E.g. “our network is fine, there are no complaints”
 - Backbone Network
 - Will assure “internal” performance, but Backbone responsibilities normally stop at the demarcation point, blame is shifted to other networks up and down stream

Motivation – A Typical Scenario

- Stumbling Blocks to solving performance problems
 - Lack of a clear process
 - Knowledge of the proper order to approach problems is paramount
 - This knowledge is not just for end users – also for application developers and network operators too
 - Impatience
 - Everyone is impatient, from the user who wants things to work to the network staff and application developers who do not want to hear complaints
 - Information Void
 - Lack of a clear location that describes symptoms and steps that can be taken to mitigate risks and solve problems
 - Lack of available performance information, e.g the current status of a given network in a public and easily accessible forum
 - Communication
 - Finding whom to contact to report problems or get help in debugging is frustrating

Motivation – Possible Solutions

- Finding a solution to network performance problems can be broken into two distinct steps:
 - Use of *Diagnostic Tools* to locate problems
 - Tools that actively measure performance (e.g. Latency, Available Bandwidth)
 - Tools that passively observe performance (e.g. error counters)
 - *Regular Monitoring* to establish performance baselines and alert when expectation drops.
 - Using diagnostic tools in a structured manner
 - Visualizations and alarms to analyze the collected data
- Incorporation of either of these techniques must be:
 - *ubiquitous*, e.g. the solution works best when it is available everywhere
 - *seamless* (e.g. *federated*) in presenting information from different resources and domains

Diagnosis Methodology

- Find a measurement server “near me”
 - Why is this important?
 - How hard is this to do?
- Encourage user to participate in diagnosis procedures
- Detect and report common faults in a manner that can be shared with admins/NOC
 - ‘Proof’ goes a long way
- Provide a mechanism for admins to review test results
- Provide feedback to user to ensure problems are resolved

Partial Path Decomposition

- Networking is increasingly:
 - Cross domain
 - Large scale
 - Data intensive
- Identification of the end-to-end path is key (must solve the problem end to end...)
- Discover measurement nodes that are “near” this path
- Provide proper authentication or receive limited authority to run tests
 - No more conference calls between 5 networks, in the middle of the night
- Initiate tests between various nodes
- Retrieve and store test data for further analysis

Systematic Troubleshooting Procedures

- Having tools deployed (along the entire path) to enable adequate troubleshooting
- Getting end-users involved in the testing
- Combining output from multiple tools to understand problem
 - Correlating diverse data sets – only way to understand complex problems.
- Ensuring that results are adequately documented for later review

On Demand vs Scheduled Testing

- On-Demand testing can help solve existing problems once they occur
- Regular performance monitoring can quickly identify and locate problems before users complain
 - Alarms
 - Anomaly detection
- Testing and measuring performance increases the value of the network to all participants

Our Goals

- To spread the word that today's networks really can, do, and will support demanding applications
 - **Science**
 - Physics
 - Astronomy
 - Biology and Climate
 - Arts and Humanities
 - Computational and Network Research
- To increase the number of test points
 - Instrumenting the end to end path is key
 - Spread the knowledge and encourage adoption

Other Thoughts

- See a talk from the recent Joint Techs Conference:
 - <http://www.internet2.edu/presentations/jt2010july/20100714-metzger-whatnext.pdf>
- Take home points:
 - Close to \$1 Billion USD spent on networking at all levels (Campus, Regional, Backbone) in the next 2 years due to ARRA Funding
 - Unprecedented access and capacity for many people
 - Ideal View:
 - Changes will be seamless
 - Completed on time
 - Bandwidth will solve all performance problems
 - Realistic View:
 - Network 'breaks' when it is touched (e.g. new equipment, configs)
 - Optimization will not be done in a global fashion (e.g. backbone fixes performance, but what about regional and campus?)
 - Bandwidth means nothing when you have a serious performance problem



Welcome & Performance Primer

August 10th 2010, OSG Site Admin Workshop – Network Performance
Jason Zurawski – Internet2

For more information, visit <http://www.internet2.edu/workshops/npw>