



Deploying a Statewide Performance Measurement Infrastructure in North Carolina

Joint Techs – Summer 2009 @ Indianapolis

Tom Throckmorton <throcm@mcnc.org>

Connecting North Carolina's Future Today

Why do we need statewide network measurement?

- NC School Connectivity Initiative
 - Over the past year, 115 school districts connected to NCREN
- We were already the primary means of connection to I2 R&E / CPS, commodity for the majority of our public and private higher ed. institutions
- Altogether, ~240 “customers”
- Emphasis on network performance, not just network availability or usage

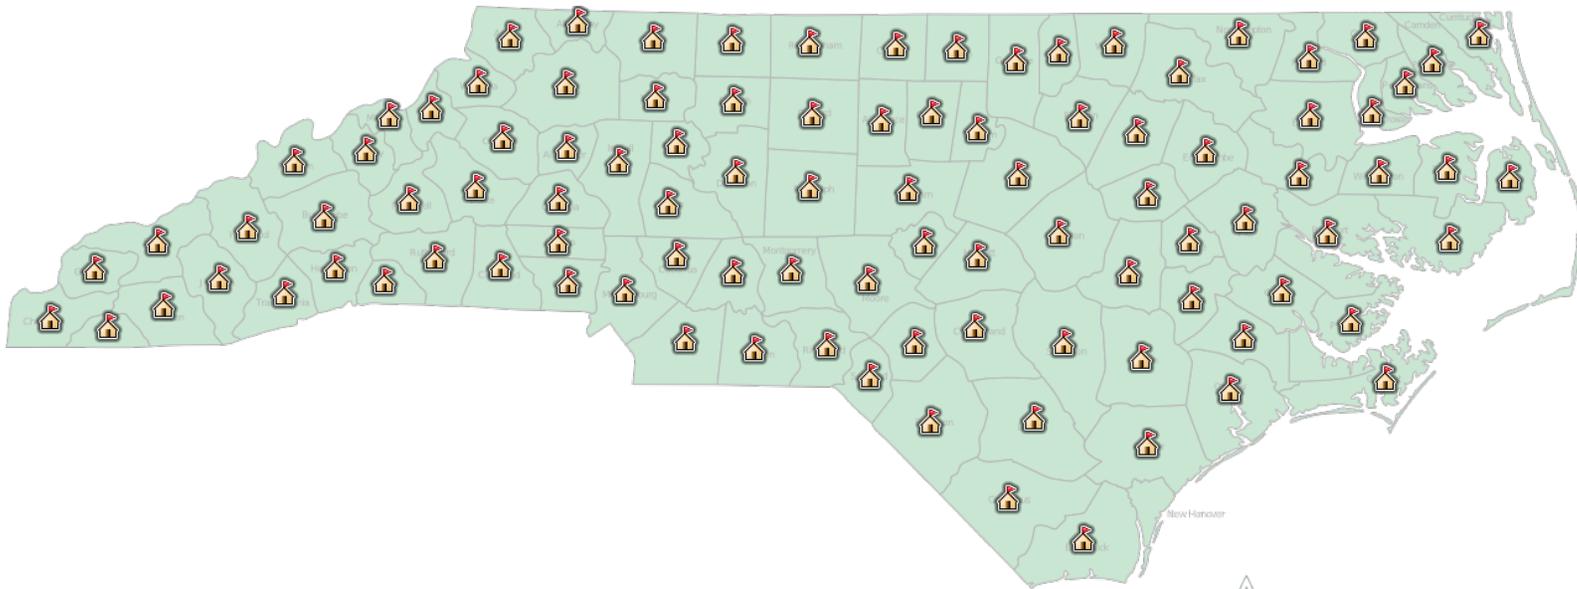


K20 Network Performance Vision

- Goals for network performance measurement
 - K20 community, all-inclusive
 - non-intrusive
 - comprehensive, transparent view of the network (strategic initiative)
 - network & application characterization
- where we want to be
 - increase active measurement on the backbone
 - continue measurement out to the customer edge
 - distributed, maintainable & extensible (open) tools foundation

Challenges, technical and otherwise

- A large number of distributed collection points
 - Trying to accurately represent measurement relevant from the perspective of our constituents, not from ours
 - Providing meaningful and relevant output, visualization
 - considering all active measurement points, collection and visualization can get complex quickly...



Challenges (cont.)

- centralized collection of the data
 - easy when you're collecting passive data (counters, etc.) or want the perspective from a central point, otherwise...not
- Few tools - mostly geared toward non-distributed use/collection, or require instrumented endpoints
- Keeping the moving bits simple, maintainable
 - deployment and management of remote systems
 - not making sysadmins out of network engineers – or vice versa
 - Use, care and feeding not reliant on experts
- PoP's are typically DC powered – limited hardware selection
- ...and limited space ☺

problems we're trying to solve

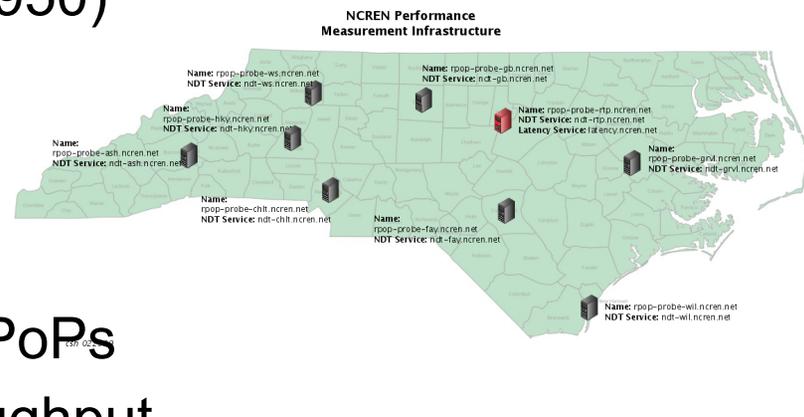
- Constituent connection issues
 - “My connection seems slow / was fast, is now slow / is not as good my connection to X”
 - No “typical” connection on the customer side – firewalls, packet shapers and bears (oh, my!)
 - Varying levels of expertise
 - Customer network documentation? Bueller?
- Supporting SLA targets

Measurements and tools

- Things we were already doing
 - some round-trip latency (from one to many)
 - Ad-hoc throughput testing (NDT)
 - Interface stats and flow collection
- Things we wanted to add / change
 - Latency measurement from end to end (and points in between)
 - Regular throughput baselining (across our backbone)
 - Push measurement closer to the users

Solution

- Measurement probes in 9* regional PoPs (plus one for test/staging)
 - commodity systems (Dell 2950)
 - Linux (CentOS 5.x)
- Multi-purpose
 - Latency measurement
 - Round-trip, between all PoPs
 - Customer on-demand throughput
 - NDT as before, but now closer to the customer's point of connection
 - Diagnostic tools for Operations / Engineering staff
 - CLI versions of latency, throughput testing



Things we've explored / tried (so far...)

- Round-trip latency, using a variety of ICMP echo probes
- One-way latency (using OWAMP)
- Baseline throughput testing (using iperf/BWCTL)
- Distributed scheduling/collection engines
 - collectd, Cacti, SmokePing master/slave, pS-PS
- Visualization
 - Graphs, matrices, maps

Active measurement

- First active measurement application - round-trip latency
 - Replaced customer-facing interface for cross-network latency
 - ICMP round-trip measurements between all probes
 - alerts on latency/loss deviations, thresholds
 - basis for ensuring SLA is met
 - new latency matrix (<http://latency.ncnrc.net/>)
 - Replaced existing Tavve eProbe infrastructure – no more recurring licensing \$\$ 😊
 - SmokePing engine today; already prototyped higher-resolution distributed collection engine (collectd)

Experience

(is what you get when you didn't get what you expected)

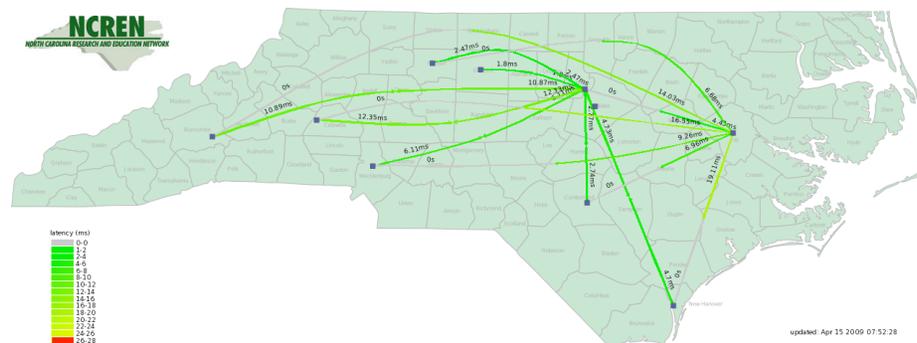
- Things we learned along the way
 - Having multiple measurement tools coexist on the same system is challenging
 - One size does not fit all (esp. cabinets)
 - Standardized, automated build and configuration management are essential to supporting remote systems
 - Found a few good tools / techniques for doing distributed monitoring / centralized collection
 - Having a server w/ some useful test tools close to the customer demarc can be handy

Advantages

- Key Advantages of new architecture
 - More active, richer measurements on the backbone
 - Distribution puts measurement closer to the constituents
 - Having views of specific areas of the network from different vantage points allows us to isolate and pinpoint problems
 - Open measurement infrastructure means potentially more measurement points, further out to the edge
 - Base and tools are open source, in wide use
 - Gives us flexibility to make local tool improvements, contribute changes back to the community
 - Collaborate, share best practices with peer institutions
 - Tools replace some existing functionality, allowed us to retire older measurement components
 - foundation for future flexibility

What's next?

- Continued consolidation / integration
 - Look at consolidating existing services onto this infrastructure
 - multicast beacons next on the list
 - Feed measurement results into availability monitoring
- Continue to seek out opportunities for using, extending this platform
- One-way, throughput measurements
 - Timing / scheduling are critical
- Visualization of collected data
 - Continue to extend geographic views of the data
 - Region-specific views
 - Representation of throughput, current network status (weathermap?)



And beyond...

- Looking beyond the PoPs
 - Looking at increasing constituent network visibility by increasing measurement to the customer edge using:
 - Low-cost NIDs for performing non-obstrusive in-service throughput measurement, latency/loss/jitter, layer 1/2 diagnostics
 - Or perhaps something locally-built (plug computers?)
 - Seek out and encourage constituent participation in measurement
- Interested in hearing feedback and/or sharing experience with anyone trying to achieve similar results

Thanks!

Questions?

Tom Throckmorton throck@mcnc.org