



NetDev 0.1

TCP Stack Instrumentation BoF

Agenda

- PLEASE SIGN IN!
- Who Am I?
- What Is Stack Instrumentation
- Use Cases
- Implementation Questions
- Where do we go from here?

Who Am I?

- Chris Ravier
 - Pittsburgh Supercomputing Center
 - Supporting scientific users from many fields
 - Provides storage and compute resources
- Primary work in network performance and diagnostics
- Associated with HPN-SSH and Web10g
 - Web10g implements RFC 4898 providing a wide range instruments

What is instrumentation

- What do we mean by instruments?
 - A method of collecting data regarding per flow characteristics not normally revealed by the stack.
 - Collected metrics must be revealed to userland
 - Provides real and quantifiable value to some set of users (engineers, diagnostics, research, etc)
 - How to validate the instruments
 - Quantification of impact on performance

Use Cases

- My community wants real world data to enable flow analysis and diagnostics
- What else? Which instruments?
 - Eg: Twitter wants to know whats happening between containers (with their own stacks) when their jobs interfere with each other. They want pkt/byte stats, congestion metrics, RTT, queue lengths.

Implementation

- How to implement the instruments
- How to provide the data to userland
 - TCP_INFO
 - Obvious choice
 - NETLINK
 - Something else?
- Do we even need to address this question now?

Where do we go from here?

- Mailing list
- Working group
- Future conference meet ups
- Coordinated effort
 - What are the shared goals?