ESnet
Energy Sciences Network

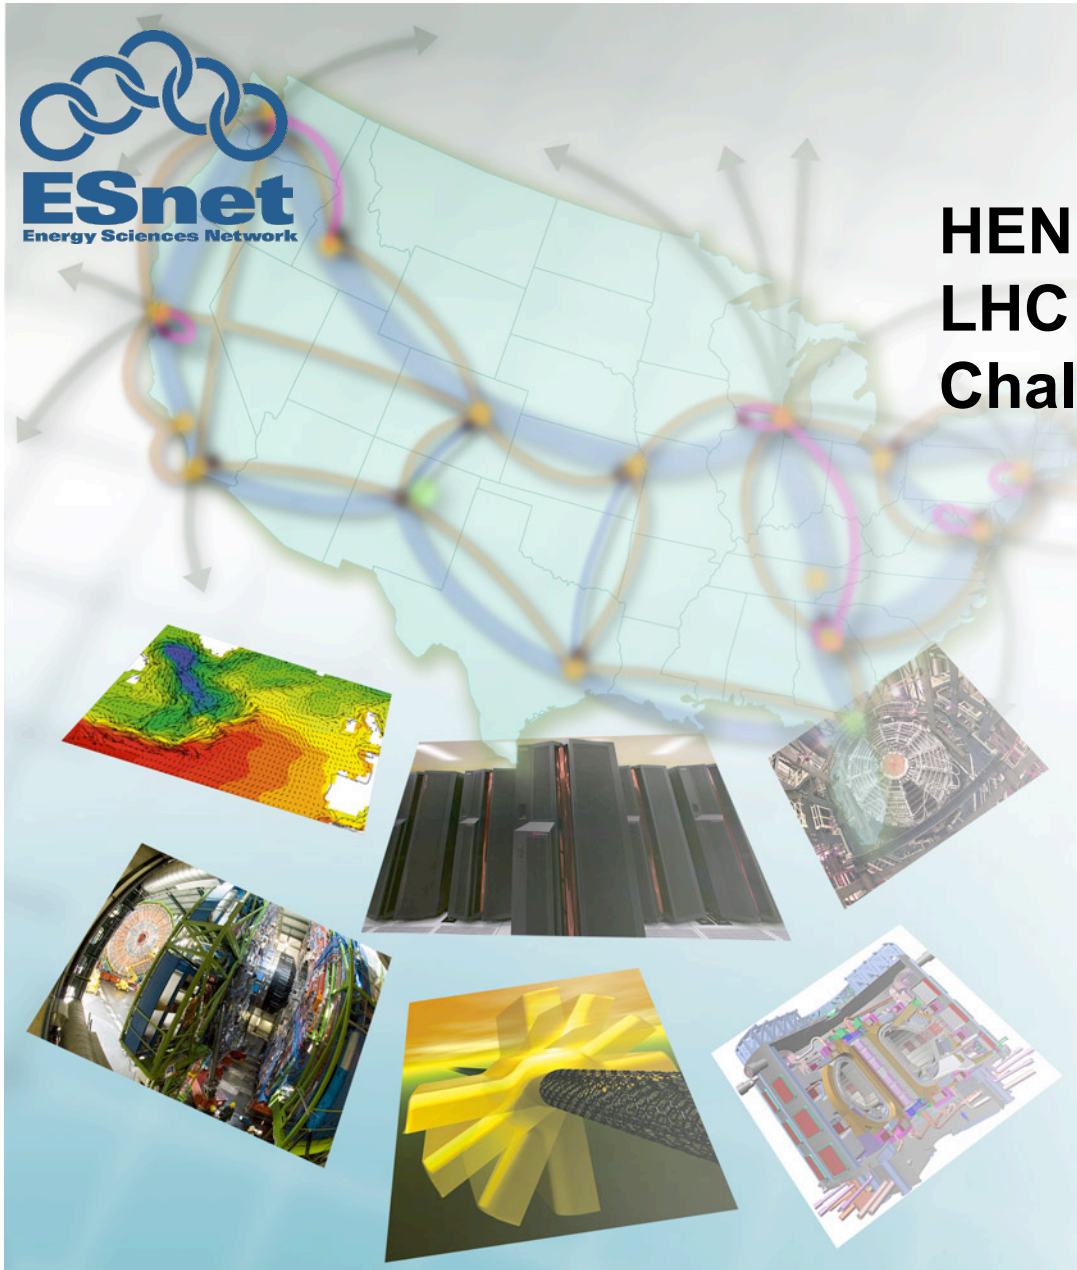

HENP SIG
LHC Network Measurement Challenges

Joe Metzger

Energy Sciences Network
Lawrence Berkeley National Laboratory

April 22 2008
Internet2 Spring Members Meeting

Networking for the Future of Science

BERKELEY LAB

Office of Science
U.S. DEPARTMENT OF ENERGY

Current Network Environment

- Most R&E network backbones are composed of 10Gbps links
- The LHC community has the tools, techniques, infrastructure & capability to transfer data at 10Gbps.
- But…
  - Network topology is **constantly** changing!
  - LHC data transfer flows are not typical internet flows
  - Many network operators don't have a lot of experience with large flows
  - Most physics flows cross multiple domains
  - Many cross-domain links haven't been tested at capacity
  - Line rate flows don't aggregate nicely
  - Debugging problems can be difficult

Measurement Requirements

- You must have the ability to easily determine the status of the set of paths you rely on for your critical missions.
  - Up and working correctly?
    - **How do you prove it?**
  - Down
    - Is there a known problem that is being worked on?
      - Are you seeing a symptom of the problem or something else?
    - Is part of the network down or the applications down?
    - How do you prove the problem is, or is not in your cluster/campus/ regional?
    - Who do you call and **what hard data can you provide to help them quickly identify the problem and fix it?**
  - Up but not performing as expected.
    - Is there a known problem?
    - Who do you call and **what hard data can you provide to help them quickly identify the problem and fix it?**
- Do you know if your use of the network is affecting others?
  - Are you getting more, less, or exactly your fair share?

New Network Traffic Profiles

- Old Typical Traffic Patter

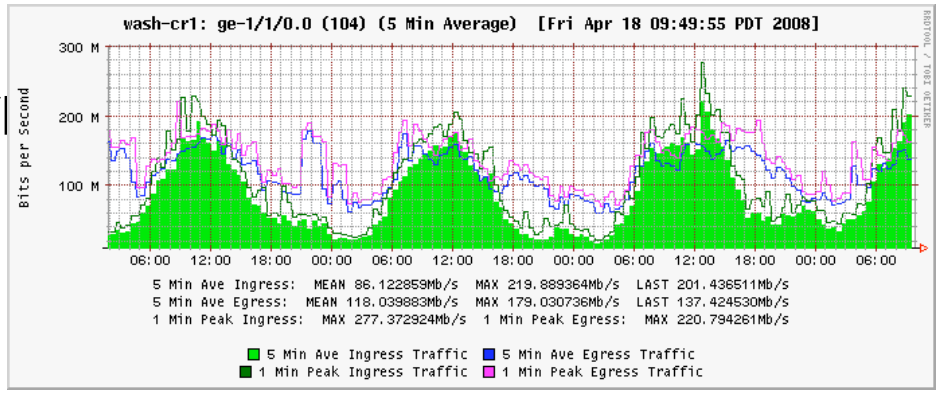

- Steady State Instrument Output Pattern

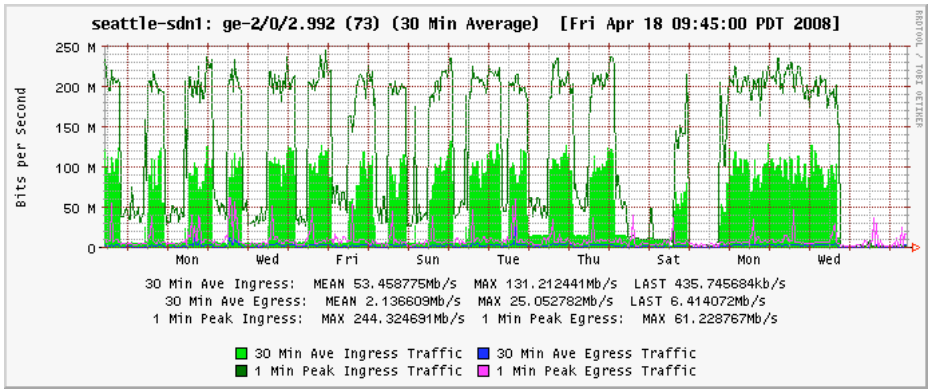

- Tuned Bulk Data Transfer

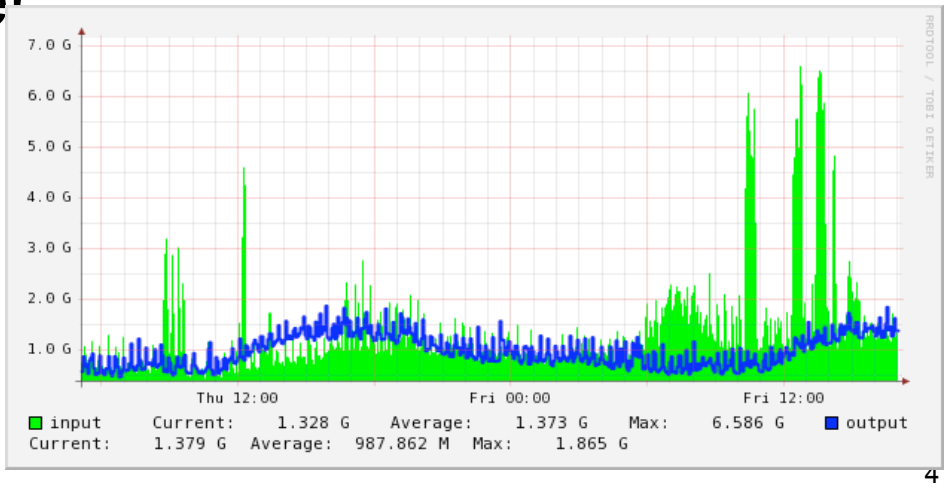

How many 5-7 Gbps flows can you aggregate on a 10 Gbps backbone?

Community Progress

- There has been a lot of work in the Network Measurement space
  - Developing frameworks for exchanging measurement data
  - Developing & improving measurement tools
  - Defining diagnostic methodologies
  - Analysis techniques
- There is a small community that understands how to use these tools and techniques for network performance analysis, verification and debugging
- We need to help the LHC community take advantage of these capabilities
  - Metcalfe's law - The value of a network is proportional to the square of the number of users.

LHC US Tier 1/2/3 Measurement Documents

- Why
  - A general white paper describing the value of network measurement to the community.
- What
  - A BCP styled after IETF BCP 15 that covers
    - What Measurements to support
      - Delay, Bandwidth, Interface Utilization, Errors & Discards, etc
    - Protocols
      - For measurement collection: ICMP, OWAMP, iperf etc.
      - For measurement Publication & Sharing - perfSONAR
    - Schedules & parameters
      - For regularly scheduled tests
    - Data sharing guidelines
- How
  - An implementation guide describing
    - What tools to use
    - How to configure them
- Constrained Scope
  - Limit the scope to the US LHC community

Next Steps

- **Identify Tier 2/Tier 3 representatives who can participate in writing & reviewing BCP**
  - **Physics User**
  - **Campus Networking Person**
- Present draft recommendations to the US LHC community at the T1/T2/T3 meeting at BNL in May
- Evaluate the ‘US Recommendations’ applicability to the global environment at LHCOPN meeting in June
- Present recommendations & pilot implementations at Joint Techs in July
- US LHC community using infrastructure by end of summer

Conclusions

- The Physics community
  - is The premier network user at this time
  - wants to be good network citizens
- The network should not be a black box.
-
-
-