

Distributing weather data via multipoint layer-2 paths using DYNES

Malathi Veeraraghavan
University of Virginia (UVA)
mvee@virginia.edu

Steve Emmerson
Univ. Corp. of Atm. Research (UCAR)
emmerson@ucar.edu

Steve Decker and Ivan Seskar
Rutgers Univ.

Jerry Robaidek and Dale Carder
Univ. of Wisconsin

April 30, 2014
CC-NIE PI Meeting




Supported by NSF CC-NIE grant: ACI-1340910

Acknowledgment

- UVA graduate and UG students: Jie Li and Yicheng Liang
- UCAR: Mohan Ramamurthy and Pete Siemsen
- UNH: Bob Russell, co-author on VCMTP
- Internet2: Brian Cashman & Eric Boyd
- V. Tech: John Lawson and Jeff Crowder
- UVA: Ron Withers
- MAX: Tom Lehman
- IU: Ezra Kissel and Martin Swany
- RENCII: Chris Heermann
- Texas Tech: Alan Sill
- CSU: Greg Redder
- BU: Weylin Piegorsch & David Starobinski
- U. Utah (Emulab), UNM PRObE, and ESnet 100G testbed support



Wide-area distributed instrument: DYNES



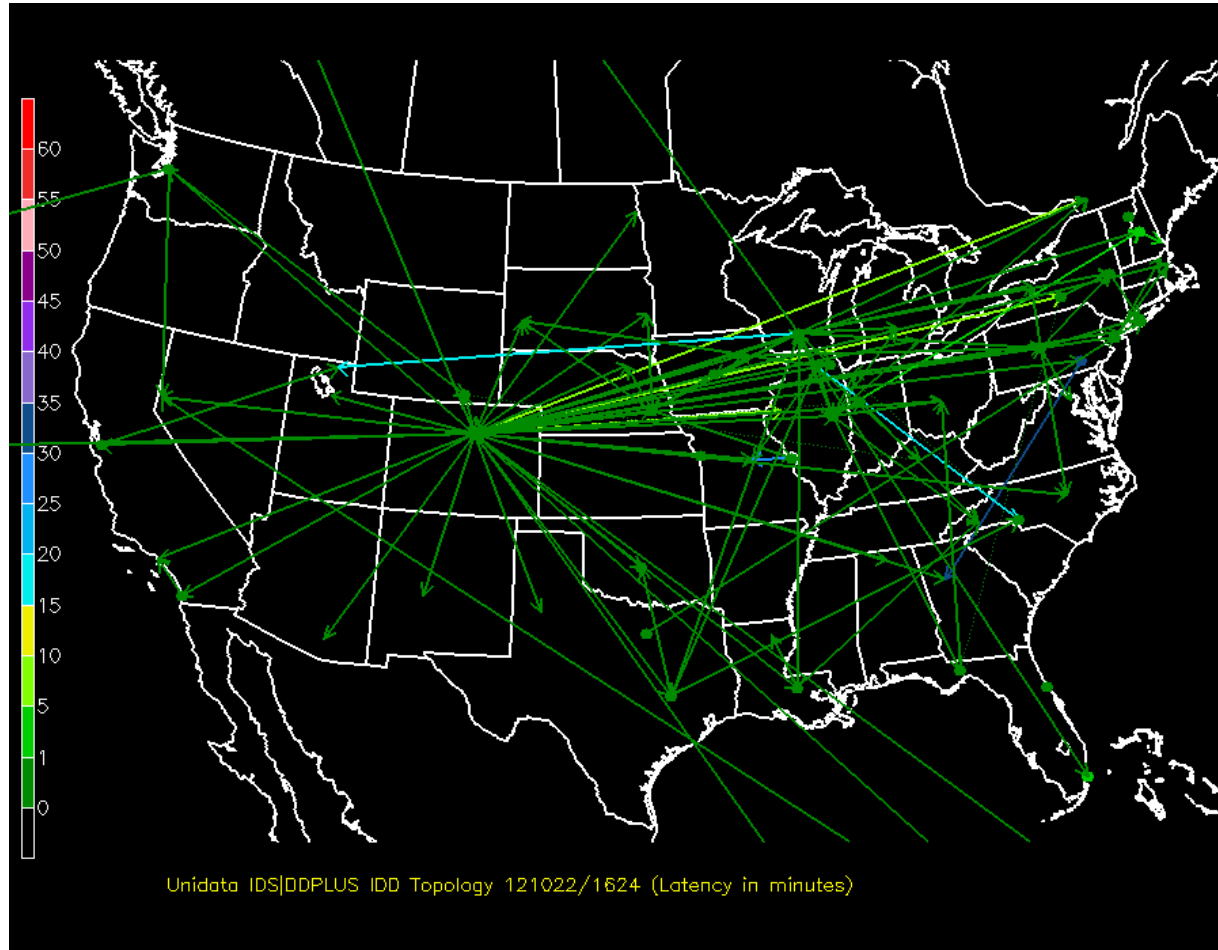
- Internet2 AL2S and ION services
- Regional services: DYNES or stitched VLANs, e.g., MARIA, FRGP

Local Data Manager (LDM)

- Software used by the Unidata Internet Data Distribution (IDD) system to deliver meteorology-related data to at least 450 computers at 240 sites (mostly universities)
- “Push” (not “pull”) of routinely-generated data-products
- 30 data feeds provide radar, satellite, text bulletins, lightning, model forecasts, surface and upper air observations, ...
- **Receives over 14 GB/h**, 24/7 (usage is increasing).
- **Transmits over 650 GB/h**, 24/7, with 99.999% reliability (usage is increasing).
- Uses RPC over **unicast TCP connections**
- Code: <http://www.unidata.ucar.edu/software/lDM/>



IDD in the USA



Global Internet Data Distribution

Unidata IDD

North American data delivery and sharing network

IDD-Brasil


South American peer of North American IDD

IDD-Caribe (planning)

Central American peer of North American IDD

Antarctic-IDD

Support of US Antarctic research community



Participants

United States
Canada
Puerto Rico
Costa Rica
Barbados
Venezuela
Chile
Brazil
Argentina
England
Portugal
Spain
Austria
Russia
Vietnam
China (Hong Kong)
South Korea
Antarctica
(incipient)



Layer-2 vs. IP vs. AL multicast

- Compared to AL multicast
 - network switches/routers making copies should (ideally) result in lower CPU and bandwidth resources for data distribution (for same latency)
- L2 virtual circuits vs IP datagrams
 - Problems with IP multicast
 - complexity of IP multicast routing
 - receivers without credentials joining
 - congestion related packet loss
 - L2 virtual circuits have setup phase and rate guarantees
- Issues with L2 paths
 - lack of ubiquitous deployment (DYNES)
 - VC setup delay (but our app. has continuous traffic)



Virtual Circuit Multicast Transport Protocol (VCMTP)

- Reliable multicast transport protocol designed for the type of data distributed by LDM
 - Almost continuous streams of small files
 - Hundreds of receivers not millions
- VCMTP design choices
 - Reliable multicast over Layer-2 virtual circuits
 - Retransmission requests/retx over TCP
 - Tradeoff throughput of fast receivers with robustness of slow receivers
 - Knob: per-file retransmission timeout factor
 - Asynchronous and message-based API
 - Receiver: per-file notification or batched notification
 - Code: <https://github.com/Unidata/vcmtip>



Project goals

- Integrate VCMTP with LDM
- Compare AL, IP and Layer-2 multicast across DYNES hosts (ten planned)
- Test on UNM PRObE with 100+ receivers



Problems encountered/ some solutions

- DYNES: thanks to Brian Cashman
 - Stitching VLANs on a per-network basis
 - Dynamic control-plane software (combination of OSCARS and OESS): TBD
 - Some sites on AL2S, some on ION
 - Need L2 traceroute: used IP addr. for VLANs/obs host
- LDM-VCMTTP software integration
 - Need changes in VCMTTP code; students graduate!
- IP multicast: have had problems
- L2 multipoint: haven't tested on AL2S
- Anticipate needing QoS config. (guaranteed rate)
 - Currently not supported by OESS on AL2S?



Questions/comments?
email: mvee@virginia.edu

