

Reliable File-Stream Multicast Application Over Multipoint Software-Defined Networks (SDN)

S. Chen, X. Ji, **M. Veeraraghavan (MV)**

Univ. of Virginia

mvee@virginia.edu

S. Emmerson

Univ. Corp. for Atmospheric Research (UCAR)

emmerson@ucar.edu

Internet2 TechX Presentation

Sept. 26, 2016

Thanks to Steve Decker, Ivan Seskar and Joe Slezak (Rutgers U), Dale Carder and Jerry Robaidek (U. Wisconsin), NSF GENI, Ilya Baldin and Yufeng Xin (RENCI)

Thanks also to Brian Cashman (Internet2), Tom Lehman (MAX), Ezra Kissel (Indiana U), Ron Withers (U. Virginia), John Lawson (MARIA), GRNOC, and several regional REN providers for their support.

Thanks to NSF for grants ACI-1340910, CNS-1405171, CNS-1531065 and to DOE for grant DE-SC0011358



Outline

- Background
- Problem Statement
- Solution
- Evaluation: GENI
- LDM7 deployment: DYNES
- Conclusions



Background

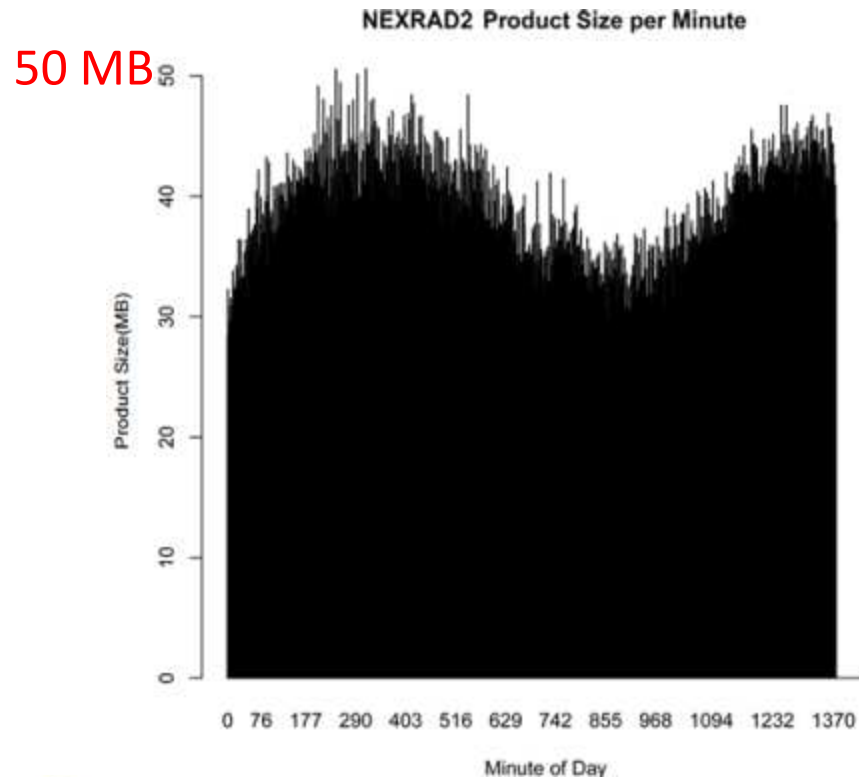
- UCAR distributes real-time meteorology data to 560 computers at **250** institutions
 - 30 data feeds: radar, satellite, text bulletins, lightning, model forecasts, surface and upper air observations
- Software: Local Data Manager v6 [since 2003]
 - Push mode; Uses **Application-Layer Multicast**
- Numbers:
 - 1 TB/hr sent by UCAR LDM servers
 - 20 GB/hr received from various sources (e.g., radar)



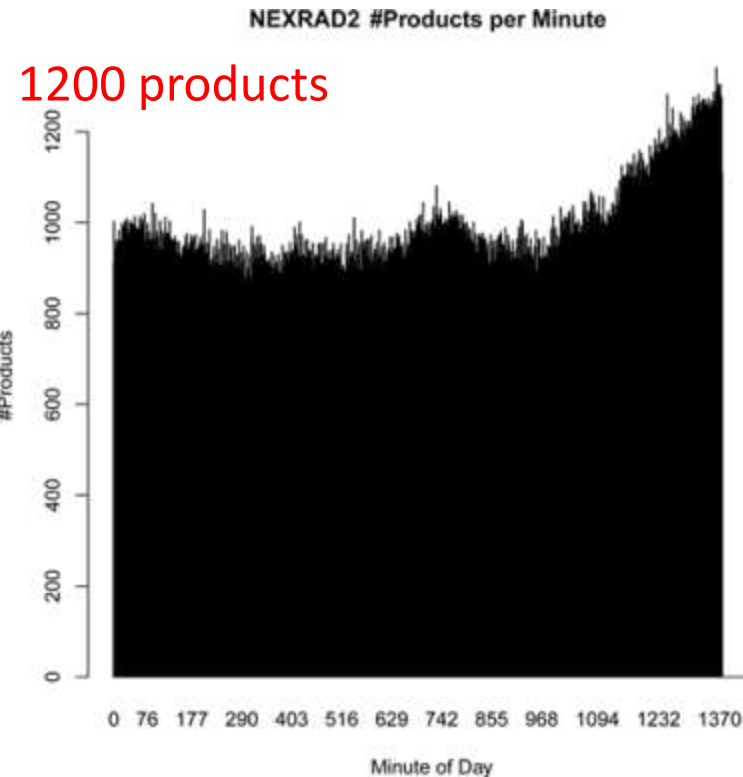
One example data feed (filestream)

NEXRAD2: Radar data

Aggregate per-minute product size



Number of products per minute

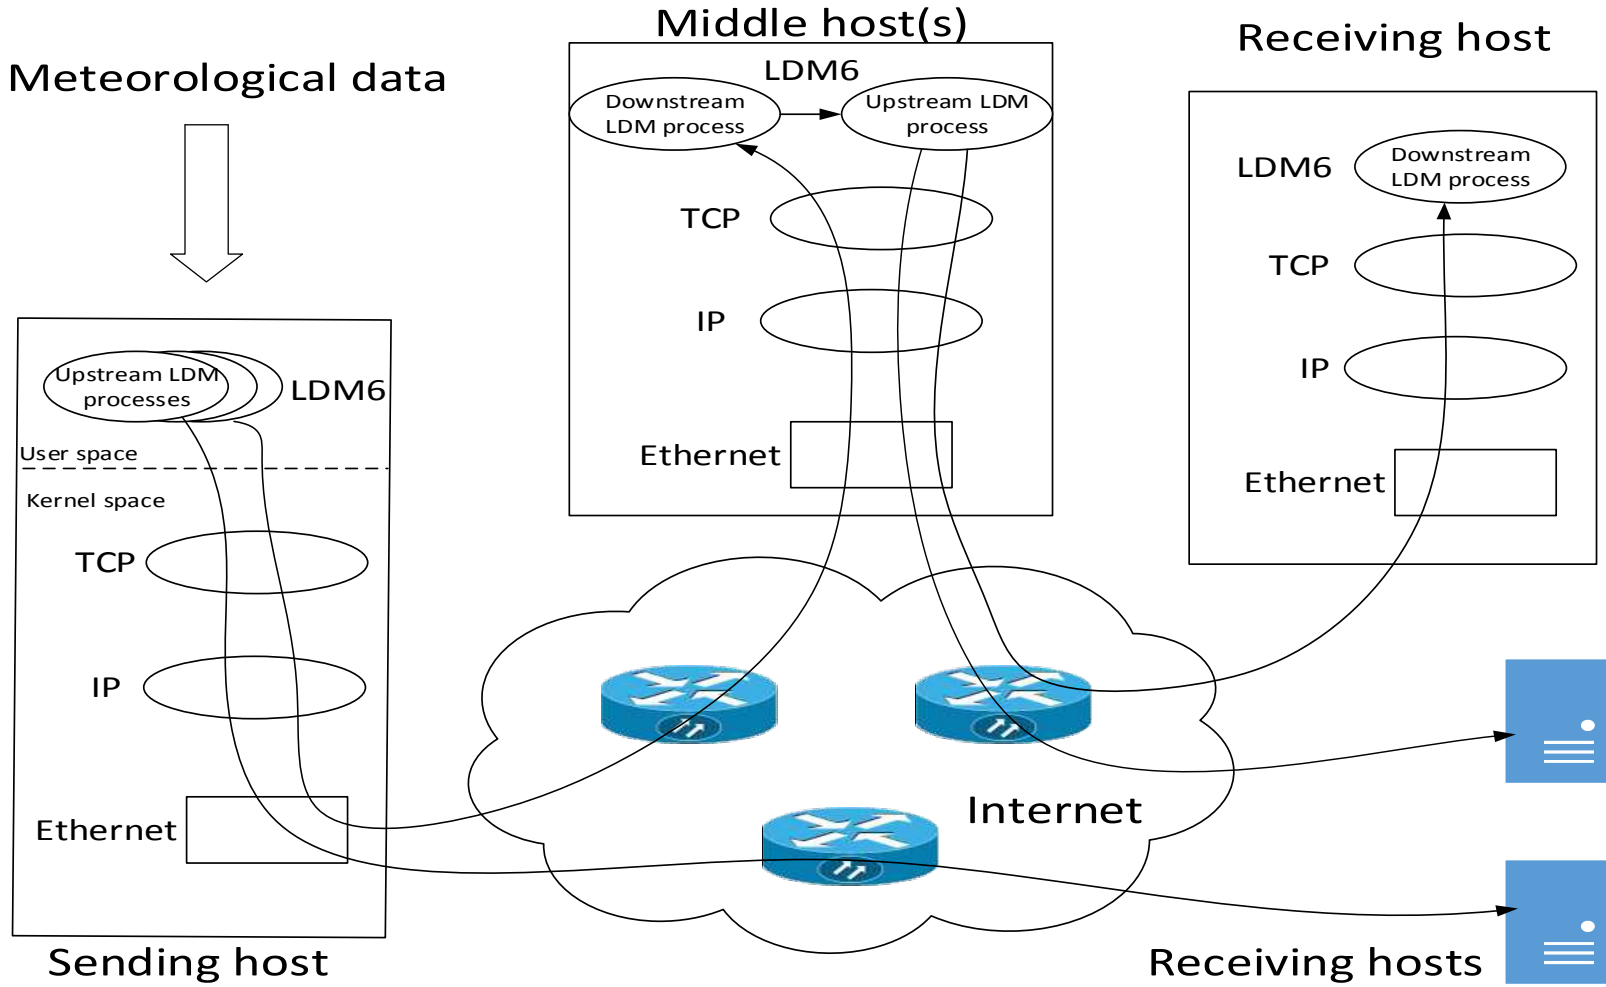


Minute of day: Traffic for 1 day for this one filestream



LDM6: current deployment

Application-layer multicast



Problem Statement

- Number of subscribers is increasing
- Volume of data is increasing
- ALM does not scale well: compute + bandwidth
- Motivation:
 - Deliver more data to more institutions
 - Deliver in less time
- LDM is in use by many institutions, not just UCAR
 - NASA, NOAA, Military, etc.
- Deployment issues with IP multicast

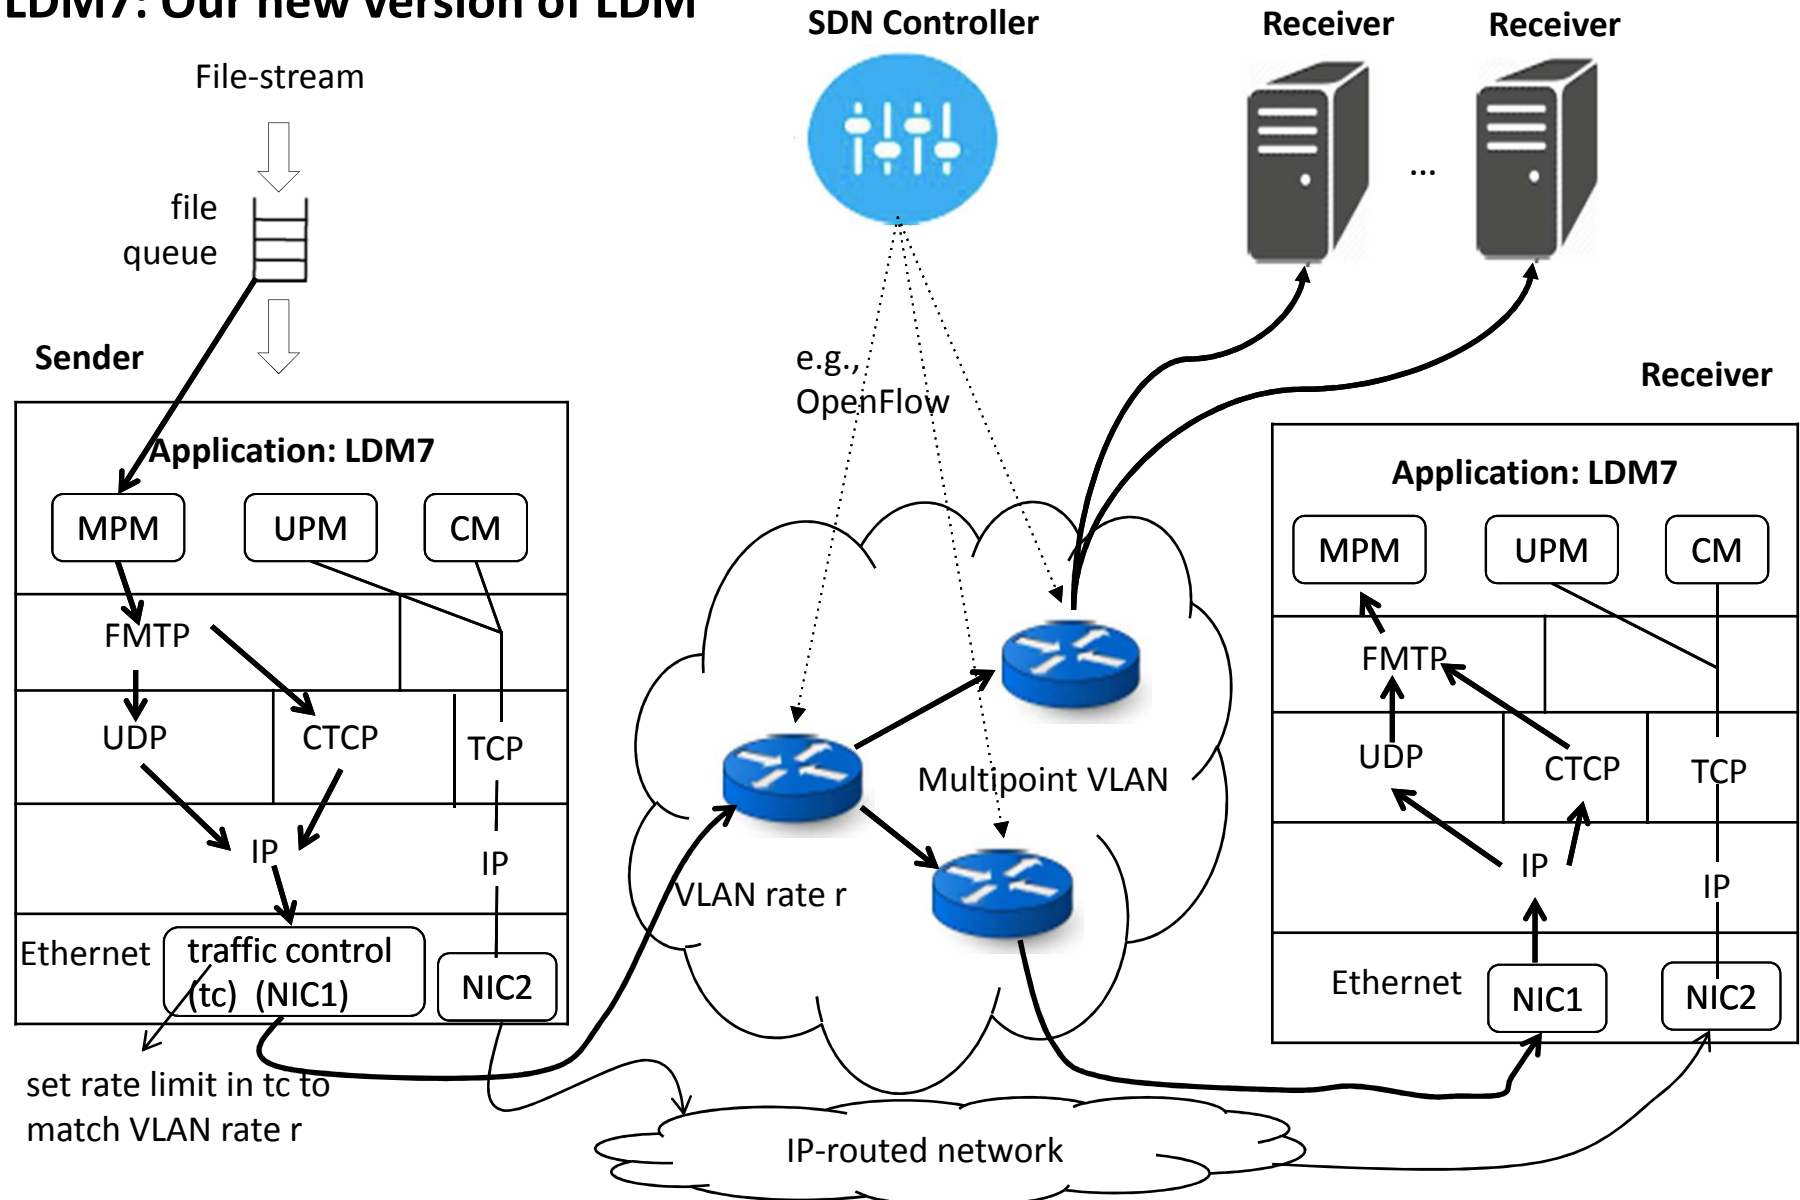


Solution: leverage SDN

- Network:
 - Set up multipoint VLAN using OpenFlow via SDN controller
 - Configure Linux traffic control (tc) in sending host
- Transport layer: File Multicast Transport protocol
 - Reliable multicast to 100s of receivers (not millions)
- Application layer - LDM7 = LDM6 + FMTP
 - API asynchronous: reduce number of copies for high-speed multicasting



LDM7: Our new version of LDM




MPM: Multicast-Push Module
 UPM: Unicast-Pull Module

CM: Control Module
 FMTP: File Multicast Transport Protocol
 CTCP: Circuit-TCP



Solution: File Multicast Transport Protocol (FMTP)

- Sender multicasts a **file (product)** segmented into **blocks (one per packet)** via UDP/IP sockets across a multipoint VLAN (L2)
- Receivers request and receive retransmissions for missed **blocks** of a file via CTCP connections; in-sequence delivery on VLANs helps detection
- Sender timer limits the service duration for a file because file held in application memory needs to be released
 - If sender receives a block retransmission request after this timer has expired, the request is rejected
 - Receiving FMTP reports unsuccessful file reception to application
 - Application receives and requests file using unicast pull module



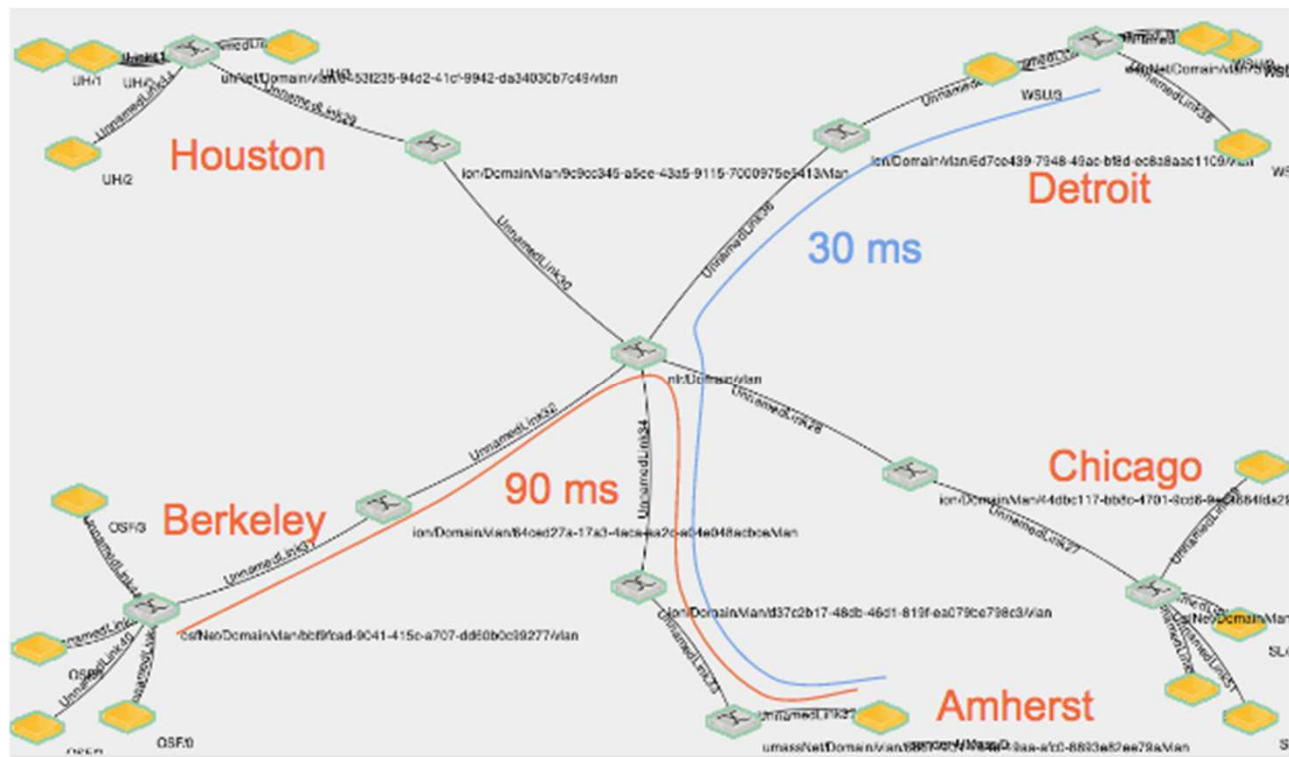
Outline

- Background
- Problem Statement
- Solution
- **Evaluation: GENI**
- LDM7 deployment: DYNES
- Conclusions



Evaluation: Experimental Setup

- 5-rack slice on NSF-funded network research testbed, GENI
- UMass (Amherst, MA), UH (Houston, TX), WSU (Detroit, MI), SL (Chicago, IL), OSF (Berkeley, CA) [RTT: 30 - 90 ms]




Evaluation metric

Average Throughput across Receivers

- LDM7: can use higher-rate VLAN (60 Mbps)
 - with $m=16$ receivers, LDM6 needs 960 Mbps
 - but still achieves lower throughput
- higher throughput \Rightarrow lower file delivery time (latency)

3-fold improvement in
Loss rate: 0% performance (file latency) Loss rate: 1%




rate of the multipoint VLAN

Evaluation metric

Resource Utilization

- Operating point for this comparison: same throughput for LDM6 and LDM7
 - LDM7 saves bandwidth (measured using sar utility)
 - LDM7 saves CPU utilization (measured using ps)

LDM7: multipoint VLAN rate = 20 Mbps; LDM6: 20 Mbps per receiver



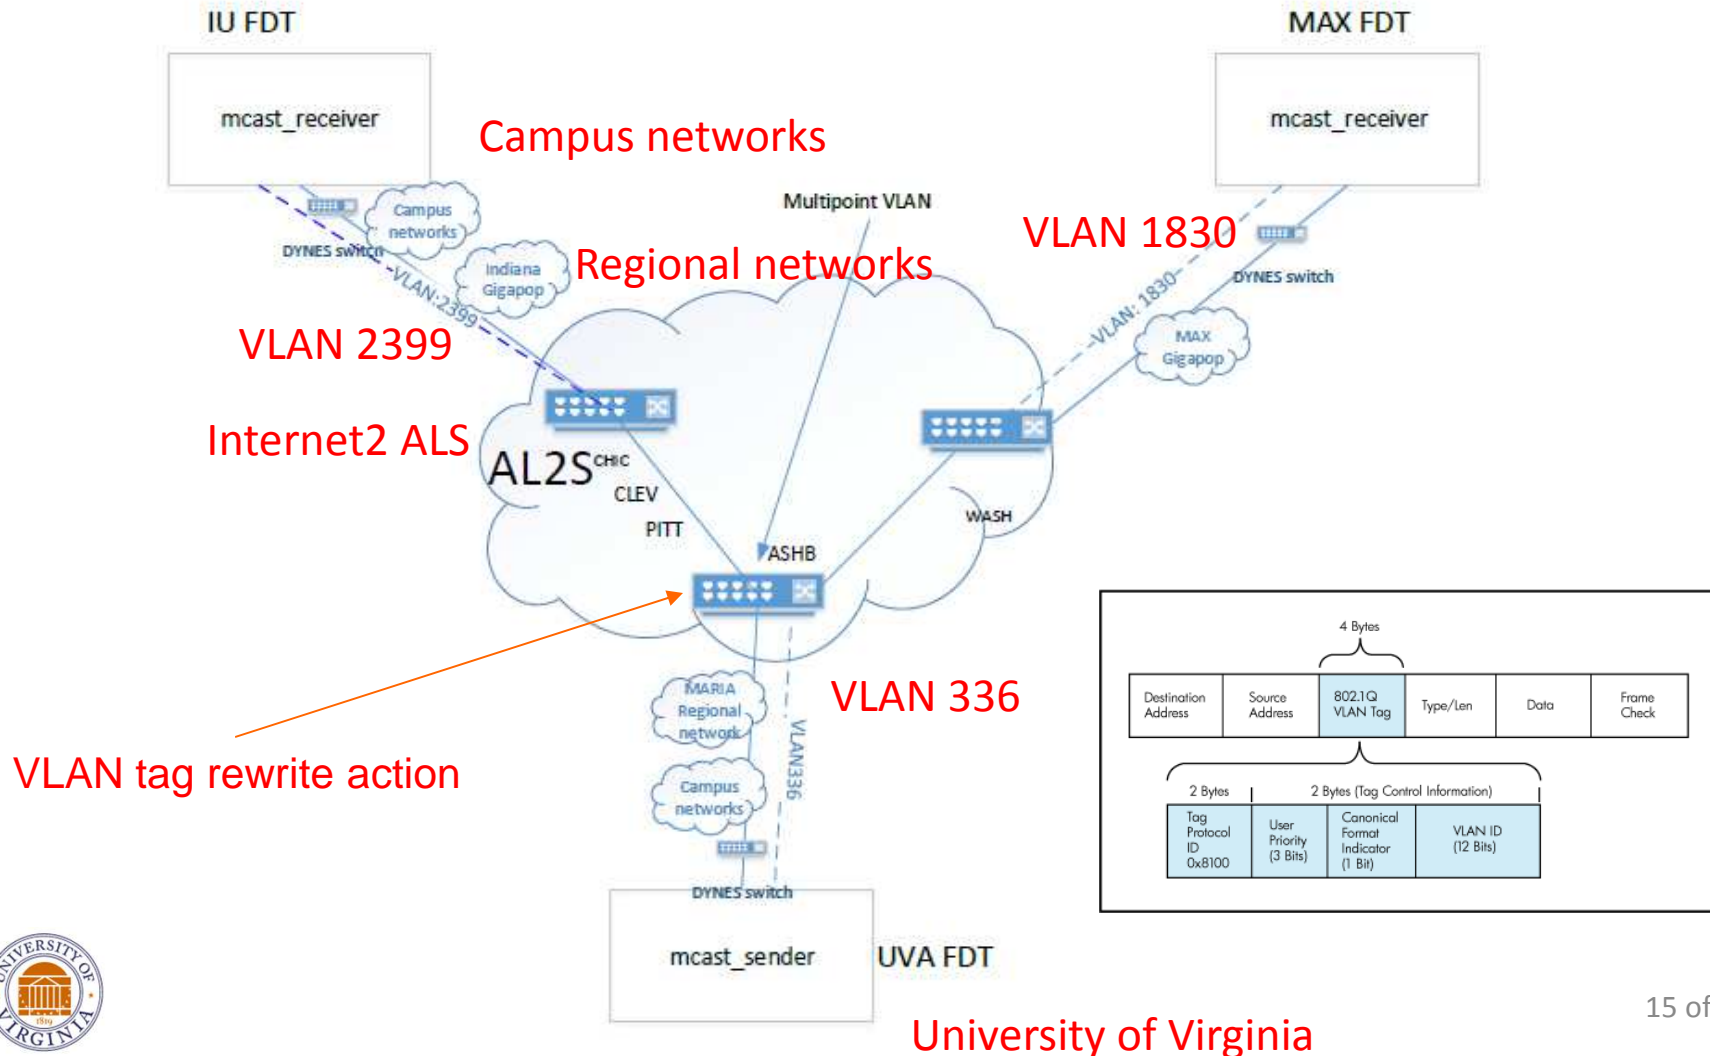
Outline

- Background
- Problem Statement
- Solution
- Evaluation
- **LDM7 deployment: DYNES**
- Conclusions



Multi-point VLAN between three university campuses via regionals and Internet2

Indiana University (Fast Data Transfer server) Mid-Atlantic Crossroads (MAX): U. Maryland



Multi-point VLAN – Control plane setup using Internet2 AL2S SDN controller

The screenshot displays the Internet2 AL2S SDN controller interface. The browser address bar shows the URL: https://al2s.net.internet2.edu/oess/index.cgi?action=view_details. The page title is "Summary".

Description: UVA-MARIA (334) to IU-IG (2399) to MAX (1830)

Summary Table:

Type	Bandwidth	Restore To Primary	Static MAC Routing	Status	Owned By
Local	0 Mbps	Off	Off	active	Virginia

Endpoints Table:

Interface	Interface Description	VLAN
sdn-sw.ashb.net.internet2.edu - et-3/0/0.0	I2-S11408 MARIA	336
sdn-sw.wash.net.internet2.edu - eth3/2	I2-S08836 MAX Gigapop	1830
sdn-sw.chic.net.internet2.edu - eth1/2	I2-S08251 Indiana UGigapop	2399

Red arrows point from the text "UVA", "MAX", and "IU" to the VLAN values 336, 1830, and 2399 respectively in the Endpoints table.

Actions: Edit Circuit, Remove Circuit, Force Reprovision, Trace Circuit Path

Utilization: History, Scheduled Events, Circuit Layout, Circuit Layout Raw

The Utilization section shows a map of the network topology and a line graph for the selected interface: sdn-sw.ashb.net.internet2.edu - et-7/0/0.0. The graph shows utilization over time, with a blue line for the main utilization and a green area for a secondary metric.



Conclusions

- SDN offers opportunity for network multicast
- Developed reliable multicast transport protocol called File Multicast Transport Protocol (FMTP)
- Integrated with long-running application: LDM
- Compared performance of current LDM6 with new LDM7 solution
 - Improved performance: lower file-delivery time
 - Improved resource utilization
- Challenge: Deployment requires VLANs in campuses

