

Optimization of Rucio-SENSE DMM



Camille Sicat, with mentorship by Aashay Arora and Diego Davila

Premise

- Current experiments, such as CMS, generate as much as 20 PB/year
- Future planned experiments, such as the the HL-LHC, will generate 30x more data

	# of collisions	# of events simulated	RAW event size [MB]	AOD event size [MB]	Total per year [PB]
Run 2	9 Billion	22 Billion	0.9	0.35	~20
HL-LHC	56 Billion	64 Billion	6.5	2	~600

The beams get "brighter" by x6
 Data taking rate goes up by x6
 Simulations go up by x3

Primary Data volume per year goes up by x30

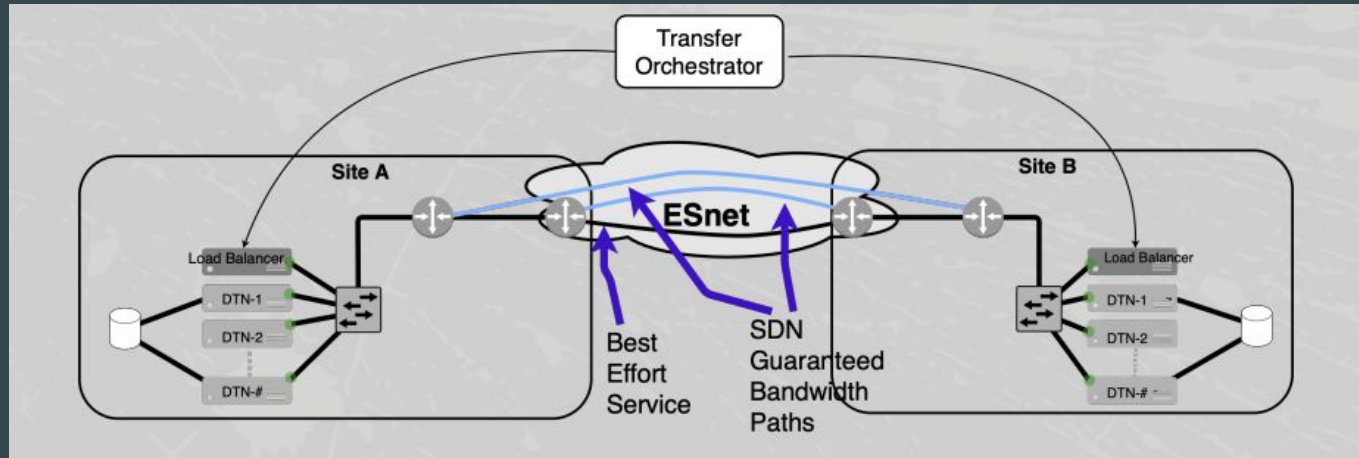
	RAW	AOD	MINI	NANO
Run 2	0.9 MB/event	0.35 MB/event	0.035 MB/event	0.001MB/event
	8 PB/year	16 PB/year	1 PB/year	0.031 PB/year
HL-LHC	6.5 MB/event	2.0 MB/event	0.250 MB/event	0.002 MB/event
	364 PB/year	240 PB/year	30 PB/year	0.24 PB/year

The Problem

- To replicate this volume of data to storage sites, need ≥ 100 Gbps network speeds
 - Estimated 400 Gbps between sites in the U.S.
- Even assuming hardware works, need to optimize network usage
- Currently, just try to push data, and if it fails, push again
 - Networks are dominated by large transfers
 - 100TB - 100PB \rightarrow 1PB @ 100 Gbps \sim 1 day
- Need more accountable network usage

Accountability with Software-Defined Networking (SDN)

- SDN allows for strong end-to-end accountability + ability to manage different priorities of workflows
- SDN integration with current tools would allow:
 - Report much data has been transferred/received
 - Comparison of expected vs. actual transfer rates
 - Identify points of failure



Rucio-SENSE DMM

- SENSE (Software-Defined Network for End-to-end Networked Science at the Exascale) allows for customizable multi-domain orchestration
 - For use on experiments with individual domain science workflows + requirements
 - Pushes QoS and routing rules
- Rucio - CERN's data management software
 - Provides scalable data storage, transfer, replication, etc. across different physical locations
 - Allows for individualized tagging/tracking of data

Current DMM Performance

- The Data Movement Manager (DMM) is the interface between SENSE and Rucio
 - From Rucio: source/destination RSE names, number of bytes, priority
 - From SENSE: Constructs P2P VLANs for each set of endpoints
 - Monitors status + performance of all dataflows
- Based on transfer metadata, DMM provisions network bandwidth for each rule
 - Constantly assesses current bandwidth usage and re-allocates depending on rule priority

Project Goals

- Make DMM more robust
 - Implementation of DMM's monitoring system
 - Correlate FTS monit data with host level information
 - Generate performance reports from combined host/transfer data (transfer rates, point of failures, etc.)
 - Optimization of DMM's interactions with Rucio
 - Optimization of DMM's bandwidth allocation algorithm


Summary

- Future experiments will generate large amounts of data
- Need to integrate SDN with existing tools to make networks more accountable for data transfers with large amounts of data
- Goal is to make Rucio-SENSE DMM more robust to meet this need

Thank you!

Backup

Detailed Model



Transfer Performance

