The LHCb Off-site HLT Farm Demonstration

Guoming Liu, Niko Neufeld
CERN, Switzerland

Motivation

The LHCb High Level Trigger (HLT) farm consists of about 1500 nodes. It is located in an underground server room (100 m underground) which suffers from:

- Limited power supply
- Limited accessibility
- Limited cooling capacity

Due to all these constraints, it is difficult to install and operate more servers in this room for the future. Off-site farm is a solution to enlarge the online computing capacity. The LHCb Data Acquisition (DAQ) is not sensitive to latency, it makes "Cloud" applicable to the LHCb HLT farm.

Background – Current LHCb Online System

The LHCb Online network is a private and almost isolated network.
- Only accessible from the CERN network via special gateway (GW) hosts.
- Use private IP subnets, not routable by CERN network.

Two dedicated parts for DAQ and ECS (Experiment Control System) for two kinds of different traffic:
- DAQ: high speed, mostly unidirectional traffic from the readout boards to the HLT farm
- ECS: low speed, bi-directional and reliable commutation for the ECS (include the HLT farm control and monitoring)

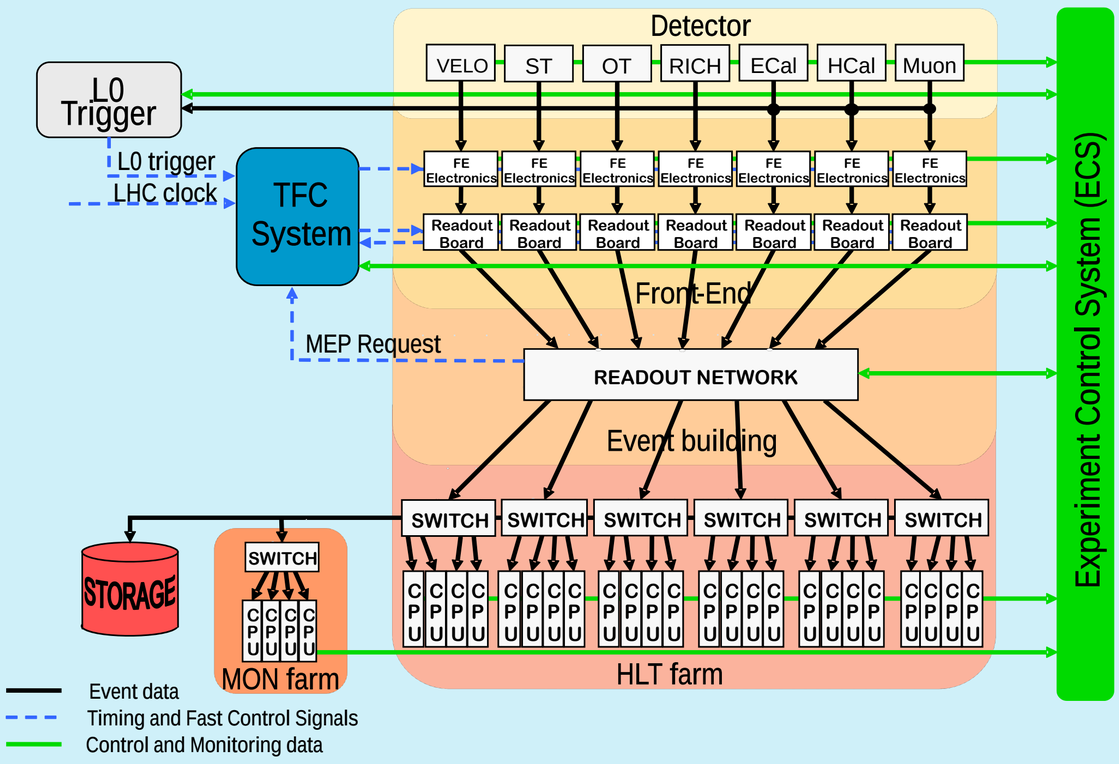

Key Technology for Off-Site Farm - VPN

A virtual private network (VPN) provides network connectivity for different sites over a possibly long physical distance.

VPN systems minimum components:
- the protocols used to tunnel the traffic
- the tunnel termination points
- the levels of security provided

Tunneling

- IP tunnel: is used to transport another network protocol by encapsulation of its packets between two networks. It may use encryption.

- NAT (Network Address Translation) :
is designed for IP address conservation which allow using private IP addresses to connect the public network. Insiders can easily see the outside world.

- Other technologies: MPLS VPN depend on support of the carrier network and were not available at CERN for this study but will be tried in the future

Off-Site Farm Demonstration Setup

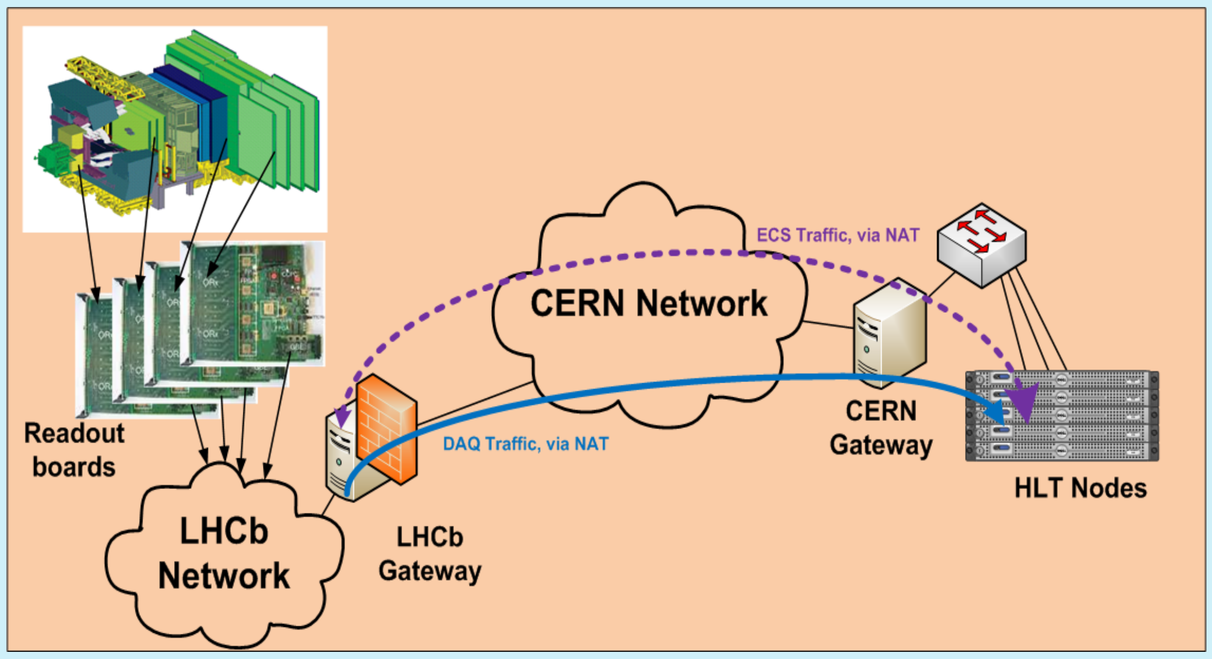

Potential solutions:
- Dedicated hardware solutions are expensive and their performance is not necessarily high enough for the DAQ traffic.
- "Software" solutions: Linux provides tunneling and encapsulating solution with great flexibility and scalability.

Our solution is a mix of IP tunneling and NAT deployed in Linux servers.
- IP2IP tunnel: used for ECS bi-directional traffic, it is good for the communication for two private sites
- NAT: used for unidirectional DAQ traffic, NAT provides high performance.

Performance Test in the Lab

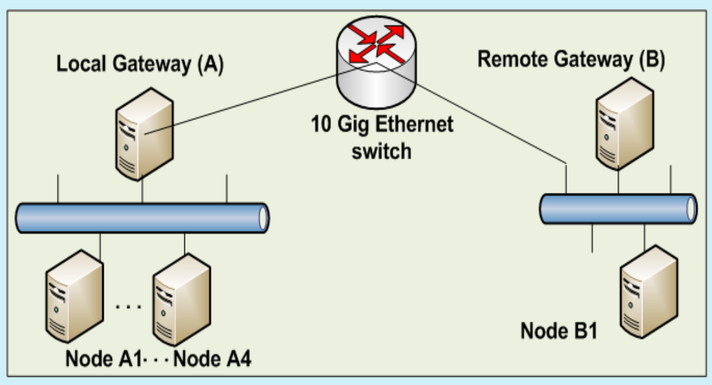

Test setup
-10Gb Ethernet
- CPU: Intel Xeon E5520
- 3 GB memory
- OS: Scientific Linux SLC 6

Test method
The local servers (A1 – A4) generate and send packets of different sizes to remote server (B1) cross IP tunneling network and NAT respectively.

Results: Throughput and CPU Load on the Gateways

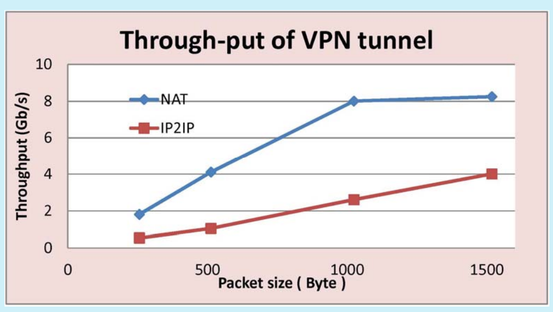

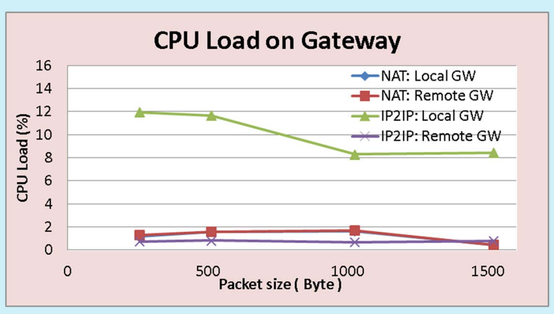

Conclusions

The concept of the off-site farm has been successfully demonstrated. A node in CERN computer center has been configured and joined the LHCb HLT farm transparently in the test, receiving test traffic but not processing. Test with real data is to be done, and the effect of network latency on the HLT performance to be further studied.

The preliminary test results on NAT and IP tunnel in the lab show good performance for the LHCb DAQ and ECS, better performance can be gained through kernel tuning and load balancing.

Guoming Liu, Niko Neufeld, CERN, Switzerland