

21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015)



Abstract ID : 291

Named Data Networking in Climate Research and HEP Applications

Abstract

Introduction

The Computing Models of the LHC experiments continue to evolve from the simple hierarchical MONARC model towards more agile models where data is exchanged among many Tier2 and Tier3 sites, relying on both strategic data placement, and an increased use of remote access with caching through CMS's AAA and ATLAS' FAX projects, for example. The challenges presented by expanding needs for CPU, storage and network capacity have pointed the way towards future more agile pervasive models that make best use of highly distributed heterogeneous resources.

In this paper, we explore the use of **Named Data Networking (NDN)** [1], a new Internet architecture focusing on content rather than the location of the data collections. As NDN has shown considerable promise in another data intensive field, Climate Science, we discuss the similarities and differences between the Climate and HEP use cases, along with specific issues HEP faces and will face during LHC Run2 and beyond, which NDN could address.

NDN

NDN, an instance of Information Centric Networking (ICN), is a new Internet architecture which focuses on the content of the data collections themselves, rather than on where the data resides. The end host addresses are replaced with content names, which, similar to URLs, are hierarchical, unique and human readable. Thus, NDN imposes minimal structure on applications, which can choose their own naming schemes. The hierarchical structure of NDN names has several advantages:

1. it is an intuitive, common organizational structure (e.g., file systems, URLs, etc.),
2. it is scalable (similar to hierarchical IP addresses), and
3. coupled with longest prefix matching, it allows for data discovery and enumeration.

NDN has a wide range of potential benefits such as in-network content caching with request deduplication to reduce congestion and improve delivery speed, simpler application configuration, and security built into the network at the data level.

The NDN concepts, structure and initial applications have been developed through an NSF Future Internet Architecture project in its second round of funding, involving eight universities. NDN has attracted significant interest from industry, including Cisco, Intel, Alcatel, Huawei, and Panasonic, and involves many of these companies through an industry consortium.

NDN and Climate Applications

We have successfully begun to test NDN in the climate application domain [2]. To handle the various naming schemes used in climate applications, we have designed and implemented translators that take existing names with arbitrary structure (produced by climate models, or home-grown) and translated them into NDN-compliant names. Depending on the original name structure, the translation can be fairly direct (e.g., data that complies with the "Data Reference Syntax" from the Coupled Model Intercomparison Project), or complex (from home-grown naming schemes that

require the analysis of metadata embedded in the dataset or even user feedback in order to construct proper NDN names).

We have deployed a dedicated 6-node testbed for climate applications that reaches locations such as the Atmospheric and Computer Science Departments at Colorado State University, LBNL and NWCS. The testbed is connected via 10G links by ESnet and is composed of high-end machines each with 40 core CPUs, 128GB RAM and 48TB disk space. The machines cumulatively host over 50TB of climate data and are used for research, experimentation and development of climate applications.

NDN Support for HEP Applications

Several features of NDN can be beneficial to the HEP computing use case. Data sources publish new content to the network following an agreed upon naming scheme. Data delivery is always performed in a pull mode, driven by the consumer issuing interest packets. Intermediate nodes in the network dynamically cache data based on content popularity, ready to satisfy subsequent interests directly from the cache, thus lowering the load on servers with popular content. Combining this with the pull-mode results in a multicast-like data delivery, possibly optimizing both the network utilisation as well as server load.

The use of multiple data sources simultaneously, as well as the native use of multiple paths between client and data source, provide for robust failover in case of network segment, node, or end-site failure.

All these are active research areas today. Caching as well as forwarding strategies, naming schemes, multi-sourcing and multi-path forwarding need to be investigated not only from the network but also the application perspective.

HEP experiments using the World-wide LHC Computing Grid (WLCG) have well-developed, hierarchical naming schemes in use, which already fit the NDN approach well. We take this logical file name structure as a starting point for investigating the benefits of using NDN as the data distribution and access network for HEP data processing. For this, we use the testbed described above. We further target simultaneous optimization of storage and bandwidth resource utilization through dynamic caching using the VIP framework in [3]. For the scalability study, we complement the testbed with the use of a simulation environment with a representative topology including network nodes and end-sites.

Summary

In this paper, we study data access over an NDN testbed developed for Climate research. We study the behaviour using HEP-like data structures based on the CMS naming scheme, showing data publishing, discovery and retrieval in an NDN network. We demonstrate the benefits of caching, speeding up data delivery in multi-job access from a single source, with jobs executing at multiple sites. We also show the results of the simulation studies of remote data access over an NDN network demonstrating the scalability of the system.

References

1. V. Jacobson, et al.; “Networking Named Content”, 2009
2. C. Olschanowsky, et al., “Supporting Climate Research using Named Data Networking”, LANMAN, 2014
3. E. Yeh, et al.; “VIP: A Framework for Joint Dynamic Forwarding and Caching in Named Data Networks”, Proc. ACM Conf. on Information-Centric Networking, 2014

Primary author(s) : SHANNIGRAHI, Susmit (Colorado State University); BARCZYK, Artur Jerzy (California Institute of Technology (US))

Co-author(s) : PAPADOPOULOS, Christos (Colorado State University); SIM, Alex (LAWRENCE BERKELEY NATIONAL LABORATORY); MONGA, Inder (ESNET); NEWMAN, Harvey (California Institute of Technology (US)); WU, John (LAWRENCE BERKELEY NATIONAL LABORATORY); YEH, Edmund (Northeastern University)

Presenter(s) : SHANNIGRAHI, Susmit (Colorado State University)

Track Classification : A07 Experiment frameworks for WAN distributed computing; A11 Computing models; T04 Data handling/access; T07 Computing facilities and infrastructures; T13 Networking
Contribution Type : oral presentation

Submitted by **SHANNIGRAHI, Susmit** on **Thursday 23 October 2014**

Last modified: Friday 24 October 2014