

Modeling Expected and Anomalous Performance of End-to-end Workflows in Extreme-scale Distributed Computing Applications

Tuesday, 11 September 2012 14:35 (12 minutes)

Abstract Purpose:

In this extended abstract, we present our broad vision of research activities that are needed to model expected and anomalous performance of end-to-end workflows in extreme-scale distributed computing applications used within the DOE communities. In addition, we present a brief summary of our current DOE-funded studies to detect and diagnose uncorrelated as well as correlated network anomaly events within PerfSONAR measurement archives collected over large-scale network topologies across multiple domains.

Application Workflow Agendas:

The next-generation of high-performance networks such as the “ANI 100Gbps network testbeds” [1] and “hybrid packet/circuit-switched network testbeds” [2] are being developed in DOE communities. They are critical for supporting research involving large-scale distributed Petascale and Exascale science experiments and their data analysis, and also cloud computing initiatives in the DOE community such as the “Magellan” [3]. They cater to the increasing network demands of distributed computing application “inherent workflow agendas” that are described in workshop reports [4] [5]. An example workflow agenda relating to bulk file transfers from research instrumentation sites can be seen in the LHC data transfers from Tier-0 to Tier-1 and Tier-2 sites. An example agenda relating to data sharing amongst worldwide collaborators for replicating results, and refining conclusions can be seen in the LHC Tier-2 site collaborations. An example agenda relating to multi-user remote instrumentation steering and visualization relates to the remote access of PNNL Confocal microscopes in GTL project. An example agenda relating to remote analytics for real-time experimentation can be seen in the ITER inter-pulse data analysis using simulation codes at remote supercomputer centers.

DOE Networking and User Expectations:

The next-generation DOE networks provide two major advantages compared to today’s networks: (i) high bandwidth capacity levels that deliver extreme-scale raw throughput performance for bulk file transfers, and (ii) very low latency levels or packet serialization times that deliver high-quality user experience for remote instrumentation steering, visualization and interactive-analytics. Consequently, the various DOE-supported distributed computing application users are generating a combination of both bandwidth-intensive and latency-sensitive traffic flows on the order of scales that have never been seen before. Given the substantial infrastructure investments to provide the high-performance networking capabilities, the networks will need to function in a manner that meets the high application performance expectations of the users. Examples of user expectations could include: (a) moving a Terabyte of LHC data within 6 hours between international collaborator sites, (b) smooth remote steering of the PNNL Confocal microscope that generates 12.5 Gbps high-definition video stream per camera to deliver “at-the-instrument” user experience for multiple geographically dispersed remote users, and (c) a west-coast remote user experiencing reliable performance over long time-periods when manipulating simulation codes and their graphical user interfaces pertaining to 2 to 3 Gbytes ITER inter-pulse data being transferred and analyzed every 20 minutes at NERSC supercomputer nodes.

Need for Novel Characterization and Modeling Strategies:

To ensure such robust functioning of next-generation networks, unique traffic flows need to be characterized and modeled to understand the user, application and network interplays. In the same context, the host and network device technologies supporting these extreme-scale applications are in their early stages of development to support 10-100Gbps application traffic flows, and will introduce performance bottlenecks that need to be detected, localized and resolved. The lessons to-be learned from such bottleneck detection testing will drive the design, development, deployment and monitoring of future 10-100Gbps, and beyond - supporting host and network technologies. Even more importantly, users/operators of the extreme-scale applications will need to have “expectation-management” tools that enable them to model, analyze and visualize if their inherent workflow agendas are performing as expected or are anomalous (particularly if they are faulty), given the infrastructure resources (e.g., instrument, compute nodes, storage, network circuit) being co-scheduled to meet their application demands. If the anomalies are benign and cross expectation boundaries, it will still be beneficial for users/operators to be notified about such changes. The expectation and change notifications could be instantaneous or could be in the form of daily or weekly trends that highlight facts such as for e.g., typically noon – 4pm on Tuesdays and Thursdays, paths of interest for the user tend to be congested due to flash crowd behaviors in recurring LHC experiments, or due to any other extreme-scale application traffic flows that increase the co-scheduling loads over the shared infrastructure of a DOE community.

Potential Characterization and Modeling Strategies:

We are envisioning research and development activities that aim to meet the above needs of extreme-scale: (i) user-and-application, as well as application-and-network interplay characterization and modeling involving bandwidth-intensive and latency-sensitive traffic flows, (ii) fault detection and localization by analyzing performance measurements across end-to-end host and network devices, and (iii) users' workflow agenda performance "expectation-management" modeling and analysis tools that extend familiar and widely-adopted middleware software interfaces (e.g., Pegasus Workflow Management System [6], Netlogger [7], perfSONAR[8], NetAlmanac [9]). These three activities should build upon each other, and our hope is that they will ultimately provide the DOE community: mathematical-models of network requirements for extreme-scale DOE user applications; fault detection and localization framework leveraging latest advances from regression analysis, model learning and constraint satisfaction theory; openly-available tools for extreme-scale application modeling and simulations, and real-network workflow agenda performance measurements. Further, we remark that existing models of application performance in the DOE community are mostly network Quality of Service (QoS) centric and focused on bulk file transfer application performance. There is a dire need to fill the dearth of knowledge in the DOE community regarding performance issues and modeling of user-and-application, as well as application-and-network interplay when considering mixtures of bandwidth-intensive and latency-sensitive traffic flows that will dominate the next-generation DOE networks.

When considering large mixtures of bandwidth-intensive and latency-sensitive traffic flows, and the nature of next-generation network paths, the measurement data volumes from sampling will be substantially large, and the user expectations of application performance will be considerably high. Consequently, there is a need to explore advanced statistical analysis coupled with effective visualization techniques for modeling user-and-application, as well as application-and-network interplay to detect bottleneck phenomena. There are already several tools that have been developed such as Pathload [10], Pathchar [11], Iperf/BWCTL [12], NDT[13], NPAD [14] to diagnose common bottlenecks such as duplex-mismatch, network congestion and hop mis-configurations along a path. However, the next-generation DOE networks will support extreme-scale distributed computing application "agendas" that have workflows involving user actions that are both bandwidth-intensive and latency-sensitive to communicate with multiple remote resources and collaborator sites. To detect bottlenecks in such application workflow agendas, novel user Quality of Experience (QoE) performance metrics and agenda-exercising tools need to be developed that are aware of user-and-application and application-and-network interplay issues, and bottleneck phenomena that may be very different from the phenomena seen in today's networks.

We believe that agenda-exercising tools that will need to be developed should be suitable for online monitoring and resource adaptation in production environments to maintain peak performance. It will be inevitable for the tools to be able to query and leverage the existence of perfSONAR measurement archives (and possibly demand the creation of new kinds of perfSONAR measurement archives) along network paths so as to reinforce analysis conclusions about network bottleneck causes. False alarms from such tools without proper reinforcement mechanisms through perfSONAR measurement archives could lead to undesirable mis-configuration of expensive resources. Hence, the agenda-exercising tools should be developed as interoperable (e.g., perfSONAR framework compliant) middleware software that can be leveraged by resource adaptation frameworks such as the ESnet OSCARS [15].

Prior Anomaly Detection and Diagnosis Research Results:

In our current research grant from DOE ASCR titled "Sampling Approaches for Multi-domain Internet Performance Measurement Infrastructures to Better Serve Network Control and Management", our preliminary results show how temporal and spatial analysis of latency and throughput performance data along with route information over multi-domain networks that were obtained using PerfSONAR web services can help in better understanding of the nature, locations and frequency of anomalous events, and their evolution over time. We are using network cliques modeling and evolution characterization techniques, and metrics (e.g., common hop to common event ratio; location affinity, event burstiness), adopted from fields such as Social Networking. We have been able to analyze "uncorrelated anomaly events" found in worldwide PerfSONAR measurement data sets, and also "correlated anomaly events" found in PerfSONAR measurement data sets between the various DOE national lab network locations. We believe our preliminary results are a major step towards modeling, end-to-end monitoring, troubleshooting and intelligent adaptations of workflows to ensure optimum user QoE in extreme-scale distributed computing applications.

References

- [1] ANI 100G Network Testbed - <https://sites.google.com/a/lbl.gov/ani-100g-network>
- [2] DOE ASCR Research – Next-generation networking for Petascale Science - <http://www.sc.doe.gov/ascr/Research/NextGen.html>
- [3] Magellan DOE Cloud Computing Initiative - <http://magellan.alcf.anl.gov>
- [4] Workshop on Advanced Networking for Distributed Petascale Science: R&D Challenges and Opportunities. April 8-9, 2008.
- [5] Workshop on Science-Driven R&D Requirements for ESnet, April 23-24, 2007.

- [6] Pegasus Workflow Management System – <http://pegasus.isi.edu>
- [7] D. Gunter, B. Tierney, B. Crowley, M. Holding, J. Lee, “NetLogger: A Toolkit for Distributed System Performance Analysis”, Proc. of IEEE MASCOTS, 2000.
- [8] A. Hanemann, J. Boote, E. Boyd, J. Durand, L. Kudarimoti, R. Lapacz, M. Swamy, S. Trocha, J. Zurawski, “PerfSONAR: A Service Oriented Architecture for Multi-Domain Network Monitoring”, Proc. of Service Oriented Computing, Springer LNCS 3826, pp. 241-254, 2005. <http://www.perfsonar.net>
- [9] ESnet Netalmanac – <http://code.google.com/p/net-almanac>
- [10] C. Dovrolis, P. Ramanathan, D. Morre, “Packet Dispersion Techniques and Capacity Estimation”, IEEE/ACM Transactions on Networking, Volume 12, Pages 963-977, December 2004.
- [11] A. Downey, “Using Pathchar to Estimate Internet Link Characteristics”, Proc. of ACM SIGCOMM, 1999.
- [12] A. Tirumala, L. Cottrell, T. Dunigan, “Measuring End-to-end Bandwidth with Iperf using Web100”, Proc. of Passive and Active Measurement Workshop, 2003 - <http://dast.nlanr.net/Projects/Iperf>
- [13] Internet2 Network Diagnostic Tool (NDT) – <http://www.internet2.edu/performance/ndt>
- [14] M. Mathis, J. Heffner, P. O’Neil, P. Siemsen, “Pathdiag: Automated TCP Diagnosis”, Proc. of Passive and Active Measurement Workshop, 2008.
- [15] ESnet OSCARS - <http://www.es.net/oscars>
- [16] C. Logg, L. Cottrell, “Experiences in Traceroute and Available Bandwidth Change Analysis”, Proc. of ACM SIGCOMM Network Troubleshooting Workshop, 2004.
- [17] R. Wolski, N. Spring, J. Hayes, “The Network Weather Service: A Distributed Resource Performance Forecasting Service for Metacomputing”, Journal of Future Generation Computer Systems, Volume 15, Pages 757-768, 1999.

Primary author: CALYAM, Prasad (The Ohio State University)

Presenter: CALYAM, Prasad (The Ohio State University)