Overflow: Multi-Site Aware Big Data Management for Scientific Work Flows on Clouds

D. Bhagyasree
M. Tech
Department of Computer Science Engineering
Bheema Institute of Technology and Science
Bellary - Adoni Rd, Dhanapuram, Andhra Pradesh

V. MALLESI
Associate Professor, HOD
Department of Computer Science Engineering
Bheema Institute of Technology and Science
Bellary - Adoni Rd, Dhanapuram, Andhra Pradesh

Abstract: The global deployment of cloud datacenters is enabling large scale scientific workflows to improve performance and deliver fast responses. This unprecedented geographical distribution of the computation is doubled by an increase in the scale of the data handled by such applications, bringing new challenges related to the efficient data management across sites. High throughput, low latencies or cost-related trade-offs are just a few concerns for both cloud providers and users when it comes to handling data across datacenters. Existing solutions are limited to cloud-provided storage, which offers low performance based on rigid cost schemes. In turn, workflow engines need to improvise substitutes, achieving performance at the cost of complex system configurations, maintenance overheads, reduced reliability and reusability. In this paper, we introduce Overflow, a uniform data management system for scientific workflows running across geographically distributed sites, aiming to reap economic benefits from this geo-diversity. Our solution is environment-aware, as it monitors and models the global cloud infrastructure, offering high and predictable data handling performance for transfer cost and time, within and across sites. Overflow proposes a set of pluggable services, grouped in a data scientist cloud kit. They provide the applications with the possibility to monitor the underlying infrastructure, to exploit smart data compression, Deduplication and geo-replication, to evaluate data management costs, to set a tradeoff between money and time, and optimize the transfer strategy accordingly. The system was validated on the Microsoft Azure cloud across its 6 EU and US datacenters. The experiments were conducted on hundreds of nodes using synthetic benchmarks and real-life bio-informatics applications (A-Brain, BLAST). The results show that our system is able to model accurately the cloud performance and to leverage this for efficient

data dissemination, being able to reduce the monetary costs and transfer time by up to three times.

1 Introduction

With their globally distributed datacenters, cloud infrastructures enable the rapid development of large scale applications. Examples of such applications running as cloud services across sites range from office collaborative tools (Microsoft Office 365, Google Drive), search engines(Bing, Google), global stock market analysis tools to entertainment services (e.g., sport events broadcasting, massively parallel games, news mining) and scientific workflows . Most of these applications are deployed on multiple sites to leverage proximity to users through content delivery networks. Besides serving the local client requests, these services need to maintain a global coherence for mining queries, maintenance or monitoring operations that require large

data movements. To enable this Big Data processing, cloud providers have set up multiple datacenters at different geographical locations. In this context, sharing, disseminating and analyzing the data sets results in frequent large-scale data movements across widely distributed sites. The targeted applications are compute intensive, for which moving the processing close to data is rather expensive (e.g., genome mapping, physics simulations), or simply needing large scale end-to-end data movements (e.g., organizations operating several datacenters and running regular backup and replication between sites, applications collecting data from remote sensors, etc.). In all cases, the cost savings (mainly computation-related) should offset the significant inter-site distance (network costs). Studies show that the inter-datacenter traffic is expected to triple in the following years. Yet, the existing cloud data management services typically lack mechanisms for dynamically coordinating transfers among different datacenters in order to achieve reasonable QoS levels and optimize the cost-performance. Being able to effectively use the underlying storage and network resources has thus become critical for wide-area data movements as well as for federated cloud settings.

This geographical distribution of computation becomes increasingly important for scientific discovery. In fact, many Big Data scientific workloads enable nowadays the partitioning of their input data. This allows to perform most of the processing independently on the data partitions across different sites and then to aggregate the results in a final phase. In some of the largest scenarios, the data sets are already partitioned for storage across multiple sites, which simplifies the task of preparing and launching a geographical-distributed processing. Among the notorious examples we recall the 40 PB/year data that is being generated by the CERN LHC. The volume overpasses single site or single institution capacity to store or process, requiring an infrastructure that spans over multiple sites. This was the case for the Higgs boson discovery, for which the processing was extended to the Google cloud infrastructure. Accelerating the process of understanding data by partitioning the computation across sites has proven effective also in other areas such as solving bio-informatics problems. Such workloads typically involve a huge number of statistical tests for asserting potential significant region of interests (e.g., links between brain regions and genes). This processing has proven to benefit greatly from a distribution across sites. Besides the need for additional compute resources, applications have to comply with several cloud providers requirements, which force them to be deployed on geographically distributed sites. For instance, in the Azure cloud, there is a limit of 300 cores allocated to a user within a datacenter, for load balancing purposes; any application requiring more compute power will be eventually distributed across several sites.

2 Existing System

•The handiest option for handling data distributed across several datacenters is to rely on the existing cloud storage services. This approach allows transferring data between arbitrary endpoints via the cloud storage and it is adopted by several systems in order to manage data movements over wide-area networks. Typically, they are not concerned by achieving high throughput, nor by potential optimizations, let alone offer the ability to support different data services (e.g., geographically distributed transfers).

•Besides storage, there are few cloud-provided services that focus on data handling. Some of them use the geographical distribution of data to reduce latencies of data transfers. Amazon's Cloud Front, for instance, uses a network of edge locations around the world to cache copy static content close to users. The goal here is different from ours: this approach is meaningful when delivering large popular objects to many end users. It lowers the latency and allows high, sustained transfer rates. The problem of scheduling data-intensive workflows in clouds assuming that files are replicated in multiple execution sites. These approaches can reduce the make span of the workflows but come at the cost and overhead of replication.

•On the other hand, end-system parallelism can be exploited to improve utilization of a single path by means of parallel streams or concurrent transfer. However, one should also consider system configuration since specific local constraints (e.g., low disk I/O speeds or over tasked CPUs) may introduce bottlenecks. One issue with all these techniques is that they cannot be ported to the clouds, since they strongly rely on the underlying network topology, unknown at the user-level.

Disadvantages:

- These existing works cannot reduce the monetary cost and transfer time.

3 Proposed System

- In this system, we propose Overflow, a fully-automated single and multi-site software system for scientific workflows data management.
- We propose an approach that optimizes the workflow data transfers on clouds by means of adaptive switching between several intra-site file transfer protocols using context information.
- We build a multi-route transfer approach across intermediate nodes of multiple datacenters, which aggregates bandwidth for efficient inter-sites transfers.
- Our proposed work can be used to support large scale workflows through an extensive set of pluggable services that estimate and optimize costs, provide insights on the environment performance and enable smart data compression, deduplication and geo-replication.

ADVANTAGES:

- Our proposed work can reduce the monetary costs and transfer time by up to three times.

4 Modules Description

Cloud Formation:

- In this module, we form the cloud. Here we generate one cloud service provider. It monitors site details, VM details, Metadata Registry & Transfer time.
- Then we generate sites. Each site has unique id then connect with CSP. It can view neighbor site details.
- Then we generate VM. Here each VM has unique id then connect with desirable site. It can view neighbor VM details.

Upload File & File Request:

- In this module, each VM can upload a file into its own storage. These details are stored in Metadata Registry.
- If another VM want to access this file, he sent the file request to Source VM.

Multipath Selection:

- In this module, the source VM wants to send the file into destination VM.
- To reduce the cost and transfer time, it must choose the shortest path between the source VM to destination VM.
- So it finds the Multipath using Multipath Selection algorithm then find the shortest path.

Smart data Compression & Replication:

- Big data size is too large. If any source VM send this big data to destination, its cost and transfer time is increased.
- To tackle this problem, we must compress this big data to small data. So we apply smart data compression technique.
- Finally, the source VM replicates its smart compressed data to destination VM.

5 Conclusion and Future Enhancements

This paper introduces Over-Flow, a data management system for scientific workflows running in large, geographically distributed and highly dynamic environments. Our system is able to effectively use the highspeed networks connecting the cloud datacenters through optimized protocol tuning and bottleneck avoidance, while remaining non-intrusive and easy to deploy. Currently, Over-Flow is used in production on the Azure Cloud, as a data management backend for the Microsoft Generic Worker workflow engine. Encouraged by these results, we plan to further investigate the impact of the metadata access on the overall workflow execution. For scientific workflows handling many small files, this can become a bottleneck, so we plan to replace the per site metadata registries with a global, hierarchical one. Furthermore, an interesting direction to explore is the closer integration between Overflow and an ongoing work [37] on handling streams of data in the cloud, as well as other data processing engines. To this end, an extension of the semantics of the API is needed.

References

[1] Azure Succesful Stories [Online]. Available: http://www. windowsazure.com/en-us/case-studies/archive/, 2015.

[2] T. Kosar, E. Arslan, B. Ross, and B. Zhang, "Storkcloud: Data transfer scheduling and optimization as a service," in Proc. 4th ACM Workshop Sci. Cloud Comput., 2013, pp. 29–36.

[3] N. Laoutaris, M. Sirivianos, X. Yang, and P. Rodriguez, "Interdatacenter bulk transfers with netstitcher," in Proc. ACM SIGCOMM Conf., 2011, pp. 74–85.

[4] Cloud computing and high-energy particle physics: How ATLAS experiment at CERN uses Google compute engine in the search for new physics at LHC [Online]. Available: https://developers. google.com/events/io/sessions/333315382, 2015.

[5] Costan, R. Tudoran, G. Antoniu, and G. Brasche, "TomusBlobs: Scalable data-intensive processing on azure clouds," Concurrency Comput.: Practice Experience, vol. 15, no. 2, pp. 26–51, 2013.

[6] R. Tudoran, A. Costan, R. R. Rad, G. Brasche, and G. Antoniu, "Adaptive file management for scientific workflows on the azure cloud," in Proc. BigData Conf., 2013, pp. 273–281.

[7] R. Tudoran, A. Costan, R. Wang, L. Boug_e, and G. Antoniu. (2014). Bridging data in the clouds: An

Environment-aware system for geographically distributed data transfers, in Proc. 14th IEEE/ACM Int. Symp. Cluster, Cloud Grid Comput. [Online]. Available: http:// hal.inria.fr/hal-00978153.

[8] H. Hiden, S. Woodman, P. Watson, and J. Ca»a, "Developing cloud applications using the E-science central platform." in Proc. Roy. Soc. A, 2012, vol. 371, pp. 52–67.

[9] K. R. Jackson, L. Ramakrishnan, K. J. Runge, and R. C. Thomas, "Seeking supernovae in the clouds: A performance study," in Proc. 19th ACM Int. Symp. High Perform. Distrib. Comput., 2010, pp. 421–429

[10] .reenberg, J. Hamilton, D. A. Maltz, and P. Patel, "The cost of a cloud: Research problems in data center networks," SIGCOMM Comput. Commun. Rev., vol. 39, no. 1, pp. 68–73, Dec. 2008.

[11] M. Isard, M. Budiu, Y. Yu, A. Birrell, and D. Fetterly, "Dryad: Distributed data-parallel programs from sequential building blocks," in Proc. 2nd ACM SIGOPS/EuroSys Eur. Conf. Comput. Syst., 2007, pp. 59–72.

[12] Y. Simmhan, C. van Ingen, G. Subramanian, and J. Li, "Bridging the gap between desktop and the cloud for escience applications," in Proc. IEEE 3rd Int. Conf. Cloud Comput., 2010, pp. 474–481.

[13] E. S. Ogasawara, J. Dias, V. Silva, F. S. Chirigati, D. de Oliveira, F. Porto, P. Valduriez, and M. Mattoso, "Chiron: A parallel engine for algebraic scientific workflows," Concurrency Comput: Practice Experience, vol. 25, no. 16, pp. 2327–2341, 2013.