CASFI Data Sharing Platform

Seoyeon Kang
sykang@an.kaist.ac.kr

Nicolas Aycirieix
ayciriei@enseirb.fr

Haewoon Kwak
haewoon@an.kaist.ac.kr

Sangman Kim
sangman@an.kaist.ac.kr

Sue Moon
sbmoon@kaist.edu

Computer Science Department, KAIST, Korea

ABSTRACT

Effective data management has become a key in promoting data sharing within measurement-based research community. In this work we propose a framework that helps us to share not only the data, but also the data management platform in order to promote collaboration across multiple institutions. We have designed and built a platform that manages measurement data and its metadata, provides basic processing primitives, and exchanges metadata with other platforms for easy local browsing of remote data. As our platform not only manages local data, but also provides consistent interface to remote data, users could browse and locate data of relevance without visiting multiple sites through different web interfaces. Any site could download our platform software package, install it, and make their metadata available to other sites with our platform.

1. INTRODUCTION

Measurement-based research benefits greatly from efficient and minimum-effort management of massive data and its metadata. As Paxson points out [8], we see the need for a "systematic approach to reduction and analysis" as soon as we amass some amount of data. Effective data management has become a key in promoting data sharing among users of a common interest.

Several archiving sites exist that post metadata of their collections, along with the data itself, if the size of the data is not prohibitively large. MAWI (Measurement and Analysis on the WIDE internet) [7] and WITS (Waikato Internet Traffic Storage) [9] manage metadata of Internet traffic data on their own web sites. CRAWDAD (A Community Resource for Archiving Wireless Data At Dartmouth) [4] archives data and tools from wireless networks. All these sites have their platforms with their own data sharing policies and mostly manage metadata manually. That is, metadata is managed in an ad-hoc fashion. DatCat (Internet Measurement Data Catalog) [6] of CAIDA (Cooperative Association for Internet Data Analysis) [2] takes a more structured approach and uses a metadata generating tool to extract necessary metadata in the XML format.

The systems listed above and other archiving sites address much of need for data management. However, we have identified additional requirements in data management. CASFI (Collect, Analyze, and Share for Future Internet) is a team of four universities who shares the same goal of measurement-driven research and need for large-scale data [3]. Members belong to different institutions and our data resides at individual member institutions. In order to share data in an easy, consistent manner, we need to share not only the data, but the data management platform. Due to the magnitude of data collected (typically, in the order of terabytes a day) and possible privacy issues, we find moving data between member institutions not feasible. We need to provide some processing capability to member institutions for close collaboration. Allman *et al.* have proposed a scalable sharing platform that "tracks the measurements themselves, but also the tools used to take and analyze the measurements" [1]. They see data processing as an integral component in data management system and discuss various approaches, such as running the analysis in a sandbox, checking the tool for being provably safe, and community-based safety rating.

In this work we propose a framework that helps us to share not only the data, but also the data sharing platform. We have designed and built a platform that manages measurement data and its metadata, provides basic processing primitives, and exchanges metadata with other platforms for easy browsing of relevant data. Our team uses a wide range of data types: from well-known packet capture formats to web crawled traces and database dumps. For data of well-known formats, we utilize existing tools to extract necessary metadata, just as in DatCat. For more free-form data, we depend on the data contributors to specify delimiters and individual fields and store the information as metadata. Sharing both data and the platform is integral to our teamwork. We offer the platform as a software package that any interested sites can download and install. Once a site identifies other sites of the same platform, then the site can include other sites in its metadata database. We believe a local integrated view of all available data to be a very convenient feature.

2. SYSTEM DESIGN

Figure 1 shows the architecture of our data sharing platform. Our platform consists of three parts: the platform software package, storage servers, and metadata database. The platform software package is the core part of our system. It has the web interface and server daemons. It is available at our web site [1]. The web interface shows the list of data files and has a search capability for users. A data contributor uploads data via web interface. Contributors' input passes to server daemon, and server daemon

PAM2009, April 1-3, 2009, Seoul, Korea.

[1] http://casfi-dsp.kaist.ac.kr

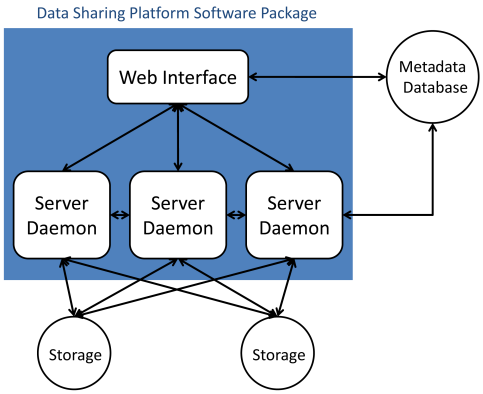

Figure 1: The architecture of CASFI data sharing platform

transforms contributors' input into meta-data in XML form. This XML format is developing with backward compatibility. Daemons pass the list of data to web interface, and web interface shows a list of data to users. Server daemons bridge web interface and storage servers. Daemons give a view of storage server to web interface, and web interface accesses data via daemons. Web interface cannot communicate with storage servers directly. This architecture improves the level of security, distributes the processing load, and offers high reliability of the entire system. The data format in all communications between the web interface and daemons follow the standard of XML-RPC. Daemons and storage servers are freely added or removed while operating data sharing platforms.

2.1 Design Considerations

In this section we consider three different roles of users; namely, a contributor, a consumer, and an administrator. A contributor uploads and shares one's own data with the public. A data consumer accesses and downloads data. An administrator installs and manages the data sharing platform. We review the needs of each role and reflect the key concerns in our platform design.

- How easy is it for a contributor to input metadata?
  We have a broad range of data from packet traces to web crawled data. We divide data into two groups based on the type: well-known form such as packet trace files and free-form. We utilize existing tools to produce metadata from well-known form data. However, for free-form data, a contributor has to specify delimiters, comment symbols, and field types. We extract the first row after comments as metadata and have users edit individual fields before depositing it to our database. We can automate the data format description process, but not for all metadata. Information, such as the router model and the link speed, cannot be inferred from the data itself, and should be maintained out-of-band.
- How can a consumer get huge or partially access-restricted data?
  Our platform provides data processing results. For well-known form data, our platform offers predefined analysis programs. We have implemented and embedded analysis programs such as analysis of size and inter-arrival time of each flow and each packet, delay of router, and utilization. For secure columns of free-form data, primitives are developed from simple aggregate functions such as `count, sum, max, min, and avg` in SQL-like fashion to advanced functions. Previous work has demonstrated the possibility of monitoring access to secure information in certain environments [5, 10]. As we have no control over the storage system at a site, we choose a different approach to address this problem. Instead of giving free reign over a data file to a user, we offer primitives. We plan to expand our set of primitives to enable more complex analysis.
- How does a contributor want to share the data?
  Our platform asks a contributor to specify primitives per column. If contributors set a specific column C, as a secure column, others cannot access full data and only can access column C with primitives. We are yet to verify the effectiveness of this fine-grain access control but, for now, we stick to our design and implementation.
- Is the data sharing platform flexible enough to reorganize the storage system?
  Modern-day storage systems come in many varieties. NAS (network-attached storage), DAS (direct-attached storage), and SAN (storage area network) are popular with computing clusters. The system administrator at a site determines how to allocate storage space to data. We propose to separate the storage management from our platform and let each site configure one's own storage system.

3. SUMMARY

In this paper, we have presented CASFI data sharing platform that manages measurement data and its metadata, provides primary processing power. We have proposed a framework to share not only the data, but also the data sharing platform that enables collaboration among multiple institutions. Since our data sharing platform is provided as a software package, it exchanges metadata with other platforms across multiple institutions in which this platform installed. For future work, we plan to download metadata from other sites in order to demonstrate ease of browsing multiple sites locally.

4. REFERENCES

[1] M. Allman, E. Blanton, and W. M. Eddy. A scalable system for sharing internet measurement. In *In Proceedings of Passive & Active Measurement (PAM)*, pages 189–191, 2002.
[2] CAIDA. http://caida.org/home.
[3] CASFI. http://casfi.kaist.ac.kr/.
[4] CRAWDAD. http://crawdad.cs.dartmouth.edu/.
[5] M. Dalton, H. Kannan, and C. Kozyrakis. Raksha: a flexible information flow architecture for software security. In *Proceedings of the 34th International Symposium on Computer Architecture (ISCA)*, pages 482–493, 2007.
[6] DATCAT. http://www.datcat.org/.
[7] MAWI. http://mawi.wide.ad.jp/mawi/.
[8] V. Paxson. http://www.icir.org/vern/talks/vp-nrdm01.ps.gz.
[9] WAND. http://www.wand.net.nz/wits/.
[10] N. Zeldovich, S. Boyd-Wickizer, E. Kohler, and D. Mazières. Making information flow explicit in histar. In *Proceedings of the 7th symposium on OSDI*, pages 263–278, Berkeley, CA, USA, 2006. USENIX Association.