IP Summarization Topic in perfSONAR Lookup Service

Marcos Portnoi
Computer and Information Sciences Department
University of Delaware
mportnoi@udel.edu

Priscilla Santos Moraes
Distributed and Meta-Systems Laboratory
University of Delaware
priscilla@damsl.cis.udel.edu

Martin Swany
Computer and Information Sciences Department
University of Delaware
swany@cis.udel.edu

Abstract

In the perfSONAR infrastructure, the Lookup Service (LS) is a key element of the measurement framework because it allows every independent service to be a visible part of the system. New services may identify themselves to the community and provide their detailed capabilities description. Other services are able to communicate to the LS in order to get this data, which is called Lookup Information. One functionality provided by the LS is IP Summarization, where subnets are aggregated by some chosen nodes. An interesting issue, which constitutes this research topic, is how to select the IP nodes where the summarization is to occur. If a node summarizes many other nodes, its load as the "keeper" might be substantial. On the other hand, if a node summarizes too few other nodes, the whole concept and advantages of aggregation are lost. This present research intends to convey a fine-tunable algorithm to select nodes for summarization, providing an efficient balance of aggregation/router load.

Categories and Subject Descriptors C.2.3 [**Computer Systems Organization**]: Computer-Communication Networks - Network operations – network management, network monitoring.

General Terms: Management, Measurement.

Keywords: IP summarization, aggregation, Radix tree, Patricia trie.

1. Introduction

perfSONAR [1] is services-oriented architecture that furnishes an infrastructure for network performance monitoring. This permits solving end-to-end performance problems on paths crossing several networks. It contains a set of services that deliver performance measurements in a federated environment. These services act as an intermediate layer, between the performance measurement tools and the diagnostic or visualization applications. This intermediate layer is aimed at making and exchanging performance measurements between networks, using well-defined protocols.

Some major services that perfSONAR provides are:

- Measurement Point Service: Creates and/or publishes monitoring information related to active and passive measurements
- Measurement Archive Service: Stores and publishes monitoring information retrieved from Measurement Point Services
- Lookup Service: Registers all participating services and their capabilities
- Authentication Service: Manages domain-level access to services via tokens
- Transformation Service: Offers custom data manipulation of existing archived measurements
- Resource Protector Service: Manages granular details regarding system resource consumption
- Topology Service: Offers topological information on networks

Particularly, the Lookup Service (LS) [2] is a key element of the measurement framework. It allows every independent service to be a visible part of the system. New services may identify themselves to the community and provide their detailed capabilities description. Other services are able to communicate to the LS in order to get this data, which is called the Lookup Information. Basically, the LS have two instances: (a) the gLS, a globally accessible Lookup Service that serves as top level of the hierarchy of LS servers, typically manage only the registration of hLS instances; and (b) the hLS, the home Lookup Service, local LS instances that manage the registration of individual services and communicate a summary of information to the upper level.

This research deals with operations a hLS must conduct to maintain the data it is responsible for. Specifically, IP summarization, which is described in more detail in the next Section.

2. IP Summarization in Lookup Service

The hLS must summarize the data set it maintains. Particularly:

- IP Addresses - The IP Addresses of all topology elements must be combined into useful CIDR style summaries. This is done through *IP Summarization*, described next.
- Host/Domain Names - Similar to IP Addresses, the hostname of each service will be extracted and broken into sub elements.
- EventTypes - The registered eventType of each data for a particular service should be extracted for future search and associated to the previous two items. The gLS will require each hLS to organize these in the summary message it registers. The eventType should be associated with both the Domain level and IP Address level summarizations.
- Keywords - If applicable, keywords can be gleaned from registered metadata and summarized in a similar style to eventTypes.

IP summarization [3] or aggregation is a mechanism used to design and maintain network data (such as routing tables) in an efficient way. Instinctively, if no method was used for route summarization, every router in a network would need to have a

route to every subnet in the network environment. This would potentially result in enormous routing tables, requiring large amount of memory and processing power from routers or specialized servers that require maintaining this sort of data.. With summarization, these servers or routers can condense some groups of routes down to a single link advertisement. Tipically, summarization is done using a special data structure called the Radix Tree or a Patricia Trie (a prefix tree) [4].

The question is: where should a network graph or tree be summarized? In other words, which IP node should be selected for summarization (all nodes below this one in the tree will not be seen by the network above this summarization node; this selected node (i.e., router) will be responsible for "knowing" the subnets below)? If a selected node summarizes too many nodes, this node will probably need to deal with a heavy load of summarization data. If, on the other hand, a node summarizes few nodes, the whole objective and advantages of aggregation are lost.

This research deals with this topic, and focuses on conveying an algorithm to decide summarization nodes at efficient points. It also intends to be tunable, so the granularity of the summarization can be controlled by an administrator.

Acknowledgments

We acknowledge the support of the Internet 2 [5] consortium.

References

[1] http://www.perfsonar.net

[2] http://wiki.perfsonar.net/jra1-wiki/index.php/Lookup_Service_resources

[3] http://www.informit.com/articles/article.aspx?p=102351&seqNum=4

[4] http://en.wikipedia.org/wiki/Radix_tree

[5] http://www.internet2.edu