

Network Working Group
Request for Comments: 3052
Category: Informational

M. Eder
Nokia
S. Nag
January 2001

Service Management Architectures Issues and Review

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

Many of the support functions necessary to exploit the mechanisms by which differing levels of service can be provided are limited in scope and a complete framework is non-existent. Various efforts at such a framework have received a great deal of attention and represent a historical shift in scope for many of the organizations looking to address this problem. The purpose of this document is to explore the problems of defining a Service management framework and to examine some of the issues that still need to be resolved.

1. Introduction

Efforts to provide mechanisms to distinguish the priority given to one set of packets, or flows, relative to another are well underway and in many modern IP networks, best effort service will be just one of the many services being offered by the network as opposed to it being the only service provided. Unfortunately, many of the support functions necessary to exploit the mechanisms by which network level service can be provided are limited in scope and a complete framework is non-existent. Compounding the problem is the varied understanding of exactly what the scope of "service" is in an IP network. IP, in contrast to connection oriented network technologies, will not be able to limit the definition of service management simply to end to end connectivity, but will combine service management with regards to transport with the service requirements of the actual applications and how they are using the network. The phenomenal growth in data networks as well as the growth in application bandwidth usage has had the consequence that the existing methods of management are not sufficient to handle the growing demands of scale and complexity.

The network and service management issue is going to be a major problem facing the networks of the future. This realization is a significant motivating factor in various efforts within the IP community which has been traditionally reluctant to take on issues of this type [1]. The purpose of this document is to explore the problems of developing a framework for managing the network and services and to examine some of the issues that recent efforts have uncovered.

2. The Problem of Management Standards

Network and service level issues traditionally are handled in IP networks by engineering the network to provide the best service possible for a single class of service. Increasingly there is a desire that IP networks be used to carry data with specific QoS constraints. IP networks will require a tremendous amount of management information to provision, maintain, validate, and bill for these new services. The control and distribution of management information in complex communications networks is one of the most sophisticated tasks a network management framework must resolve. This is compounded by the likelihood that devices in IP networks will be varied and have differing management capabilities, ranging from complex computing and switching platforms to personal hand held devices and everything in between. Scaling and performance requirements will make the task of defining a single management framework for these networks extremely complex.

In the past standardization efforts have suggested a simplified model for management on the hypothesis that it can be extrapolated to solve complex systems. This premise has often proved to be without merit because of the difficulty of developing such a model that meets both the operators heterogeneous, multi-vendor need and network equipment vendors specific needs. At the center of efforts to devise a standard management model are attempts to develop an architecture or framework to control the management information. The same conflicting operator vs. vendor forces are present in the effort to establish a common framework architecture as are in the efforts to develop a common information model.

Network operators requirements call for a framework that will permit centralized management of the network and require the minimal resources to operate and maintain while still providing tremendous flexibility in choice of equipment and creativity of defining services [2]. Operators may be less able to support change in their Operational Support Systems (OSS) then they are in the network infrastructure because the OSS is tightly integrated into the

organizations business practices. The need for flexibility, and the other desires identified above, operators expect to have meet by having equipment vendors support open and common interfaces.

Device manufactures have a need for management that will best represent the features and capabilities of the equipment they are developing and any management solution that hinders the ability of the equipment vendors to efficiently bring innovation to the market is contrary to their objectives.

The common framework for solving the management needs of operators and equipment vendors has been based on a centralized approach with a the manager agent architecture. While providing a very straightforward approach to the problem of information management, this approach, and its variations, has not proved to scale well or allowed the flexibility required in today's modern data networks. Scaling and flexibility are especially a problem where there are many sophisticated network devices present. Methods of control must be found that work and scale at the same speeds as that of the control plane of the network itself if a major concern of the management system is with the dynamic control of traffic in a network. Increasingly it is a requirement that customers at the edge of the network be able to have access to management functionality. A centralized management approach may not provide the most convenient architecture to allow this capability.

Frameworks based on a decentralized approach to the management architecture have gained momentum in recent years, but must address the possibility of having redundant management information throughout the network. A decentralized framework may have advantages with regards to scaling and speed of operation, but information and state management becomes complex in this approach, resulting in additional complication in developing such systems.

The complexity of managing a network increases dramatically as the number of services and the number and complexity of devices in the network increases. The success of IP networks can be partially traced to the successful separation of transport control mechanisms from the complexity of service management, including billing. As the trend in IP is to allow for classes of traffic that will have both transport and service dependencies it has become apparent that many of the management problems are becoming more complex in nature and are starting to resemble those of the traditional telecom provisioned service environment. In the telecom environment no such separation exists between transport control mechanisms and service. The Telecom community has struggled for years to come up with a standard solution for the problem in national and international standardization bodies and achieved a debatable amount of industry acceptance.

Unfortunately, the hard learned lessons of how to manage the interdependencies between service and transport will be of questionable use to the IP community because of the much more limited concept of service in the telecommunications environment.

Rules based management has received much attention as a method to reduce much of the overhead and operator intervention that was necessary in traditional management systems. The potential exists that a rules-based system could reduce the rate at which management information is increasing, but given the tremendous growth in this information, the problems with the control of that information will continue to exist. Rules add additional issues to the complexity of managing a network and as such will contribute to the information control problem.

2.1. IP QoS Management

Much of the current management efforts are focused on solving control issues for IP QoS [3]. A number of open questions exist with the IP QoS architecture which will make it difficult to define a management architecture until they are resolved. These are well documented in "Next steps for the IP QoS architecture" [4], but from the management perspective warrant emphasizing.

Current IP QoS architectures have not defined if the service will be per-application or only a transport-layer option. This will have significant impact both from a control perspective and from a billing and service assurance one.

The assumption is that the routing best effort path will be used for both best effort traffic and for traffic of a different service level. In addition to those issues raised in [4], best effort path routing may not be able to identify the parameters necessary to identify routes capable of sustaining distinguished service traffic.

In any architecture where a premium service will be offered it is a strong requirement that the service be measurable and sustainable. Provisioning that service will require a coherent view of the network and not just the device management view that is currently implemented in most networks.

2.2. Promise of rules-based Management

Management standardization efforts in the IP community have so far been concerned primarily with what is commonly referred to as "element management" or "device management" [5]. Generally there is agreement as to the scope of element management. Once outside that domain efforts to divide that task along clear boundaries have proved

elusive with many of the terms being used having their roots in the telecommunications industry and as such being of potentially limited use for IP management [1]. Confusion resulting from the ambiguity associated with what functions compose management beyond those intended for the element, is compounded by the broad scope for which network and service management standards apply. Terms such as business goals, service management, and application management are not sufficiently defined to insure there will not be disagreement as to the actual scope of the management functions needed and to what extent interrelationships will exist between them.

It is within this hazy domain that much of the recent efforts in rules-based management have been proposed as a potential solution. Efforts to devise a framework for policy management is an example of one of the most popular recent activities. Proposed requirements for policy management look very much like pre-existing network management requirements [2], but specific models needed to define policy itself and related to the definition of policy to control DiffServ and RSVP based QoS are under development.

2.3. Service Management Requirements

Efforts to define the requirements for a service management system are hindered by the different needs of network operators. In an industry where much has been written about the trend towards convergence there still exist fundamental differences in the business needs of operators.

2.3.1. Enterprise

The management requirements from both the operations and the network perspective have some interesting characteristics in the enterprise environment when compared to the public network. In the enterprise end to end traffic management is implemented without the burden of complex tariff issues. Service Level Agreements, while increasing in the enterprise, do not have the same operations impact as in the public network. The high costs associated with implementing non-reputable auditing systems are usually not present. This results in a substantial reduction in the number of expressions necessary to represent a particular network's business model.

In the world of best effort service, rules-based management presents the possibility to give the IT department a tool to make the network appear to not be overloaded by prioritizing traffic. This is done by prioritizing delay sensitive traffic (Web browsing) from traffic that is not delay sensitive (Email) or by prioritizing the traffic from a particular location or source. This will, depending on the composite of an enterprise's traffic, increase the useful life of the network.

without adding additional capacity. This does not come without tradeoffs. Both the purchase and management costs associated with the system must be calculated as well as the cost of the added complexity of adding additional control information to the network.

2.3.2. Service Provider

It has for a long time been a goal of service providers to have a centralized management system. While the motivation for this is very straightforward there exist some fundamental obstacles in achieving this goal. Service providers often do not want to be tied to a single vendor and certainly do not want to be limited to only one model of any single vendors equipment. At the same time bottom line costs are of paramount importance which often result in networks not being as heterogeneous as operators would like. Centralized management implies a scalable system able to manage potentially many heterogeneous pieces of equipment. The amount of data necessary to achieve this is contrary to the scalability requirement. In response to this problem it has been attempted many times to identify the common model that represents the subset common to all devices. Unfortunately all too often this set is either too complex, increasing the cost of devices, or too limited to preclude large amounts of device specific data thus defeating the purpose. For such a management model to be successful at the service level, the services being modeled must be standardized. This is counter intuitive to the competitive model of which the service provider operates. To be successful speed to market has become a key element that differentiates one service provider from another. Constraints placed on equipment manufacturers and the management infrastructure by a centralized management system are also detrimental to this goal. While for a limited set of well defined services a central management approach is feasible, such a system can very quickly become a major contributor to the very problems it was intended to solve.

3. Network and Service Management

Currently many of the efforts to define a framework for management are described in very implementation independent terms. In actual fact the implementation of that framework directly affects for what situations the management system will be most beneficial. While many past attempts to define a common management framework have failed it may be in the area of service management that such efforts finally gain industry acceptance. It may be in the domain of service management that information models can be defined that are sufficiently specific to be useful and at that same time not have a negative impact on the equipment or service providers business needs.

This section will discuss some of the issues that need to be resolved with regards to a service management framework to meet the requirements of the modern IP network.

Some of the key concerns looking at a management system architecture include:

- The management interface and models supported
- The management system architecture
- Where and how functionality is realized

3.1. Architecture for information management

Networks will consist of network elements that have existed prior to efforts to define a standard information model, rules-based or otherwise, and elements deployed after. This problem has been addressed by some of the recent efforts in policy management. Those elements that take into account policy are termed policy aware while those that do not are termed policy unaware. The distinction being made that aware devices can interpret the policy information model or schema. These issues apply equally to other standard management information. In reality it is unlikely that any device will be fully policy aware for long, as the policy information model evolves, early devices will be only policy aware for those aspects of the model that had been defined at the time. Key to success of any management framework is ability to handle revision and evolution. A number of methods exist that provide this functionality. One is designing the information models so that it can be extended but still be practically used in their original form. A second is to provide an adaptation or proxy layer. Each has advantages and disadvantages.

Methods that attempt to extend the original model often overly constrain themselves. Where the existing model cannot be extended new branches must be formed in the model that contain core management functionality.

Adaptation methods can create performance and scalability problems and add complexity to the network by creating additional network elements. A similar situation exists if the management framework is so flexible as to allow network elements to store locally information or choose to have information stored remotely. From a device perspective, the criteria will be if the device can afford the logic based on other requirements it is designed to meet, and if the information can be retrieved in such a way as to support the performance and scalability requirements that are the subject of the information. A dichotomy exists where there will be information that for reasons of performance and scalability will be transferred directly to the network elements in some situations, and in other

situations, will exist in the management plan. IP management efforts have left the level of detail needed to define the actual location of the management information to the implementation. In a service management framework it may be necessary to achieve the desired results to supply a more complete framework along the lines of detail provided by the ITU-T telecommunications management network efforts where the interfaces and functionality across interfaces has been clearly defined.

Information will need to exist in multiple locations simultaneously in any network architecture. As the quantity and complexity of that information increases limitations quickly develop. Changes in the information may need to be propagated in close to real time, further adding to the complication.

3.1.1. Rules-based Management

A network management framework can be viewed as being divided into two essential functions. The first deals with the aspects of managing the management information while the second deals with the aspects of transferring that management information into the network. The fundamental difference between rules based management and existing network management standards is that the management information is expressed as rules that reflect a desired level of service from the network as opposed to device specific management information. Many of the information management requirements of traditional management systems still apply in a rules-based environment. The network is composed of specific devices and it is at the point where rules are conveyed as device specific management information that this form of management will encounter some of its greatest challenges. A necessary component of a solution to this problem will be a generic information model to which rules can be applied and a framework architecture for distributing rules throughout the network. The task of finding the proper generic model that is not too great a burden to implement and yet provides a level of detail sufficient to manage a network has proved to be historically extremely difficult. In many ways the degree to which rules based management will be able to solve management problems is dependent on the success of efforts to define a generic model and have it be widely implemented [1].

One concept often discussed along with policy deals with the integration of legacy devices into the policy framework. The presumption is that legacy devices would be able to participate in the policy decision by having policy information translated into the native management interface. For this to succeed a device would have to support a functionality for which policy would be specified. This would limit the usefulness of this approach to only information

logically abstracted to the native interface of the device. Given that existing standard management interfaces do not support such functionality, all such devices would need to have a proprietary interface implemented. The interface being based on the existing interface supported by the device would potentially not have the scaling capabilities needed for a policy management system. Unlike a standard network management interface, were management information can be distributed between the adaptation layer and the network element, rules based management information may not be so easily distributed.

The framework for integrating rules based management system with existing network devices is not readily apparent and further study is needed. The problem exists further when one considers that there will be early policy aware devices that may not be aware as the policy models are extended. The partially policy aware devices may represent additional architectural issues as it may not be possible to expect consistency in what aspects of policy a given devices implements if there does not exist formal sets of mandatory functionality with clear evolution paths. It is paramount if the policy management framework is going to be able to evolve to accommodate the ever-increasing number of services likely to be supported by IP networks of the future that an evolution path be built into the framework.

3.2. Policy Protocol

The need for a policy protocol is important in the context of a policy aware element that is performing a certain 'service'. It is important to note here that not all elements will be aware of all service policies related to every service at all times. Therefore it makes sense for an element to be aware of a certain service policy if that element is required for a given service at any instant in time.

With the dynamics of a network where elements and links go up and down, a notion of a 'policy protocol' may become necessary. The idea of a 'policy protocol' that runs in a multi-service network requiring multi-service policies. For example; consider two arbitrary end nodes having multiple routing paths between them. Let's then assume that a certain path carries a certain service based on some Intserv bandwidth reservation technique. Let's also then deduce that the elements along that path have some element specific policy statements that have been configured on them to support that requirement. If now at any given instance any link or any element were to be unavailable along that path, the 'policy protocol' should be initiated to automatically go and configure the same service-policies

on the elements along another routed path connecting the very same end points, so that there is no disruption in service and so that no human/operator intervention is required.

The association of policy with the policy target is an area where considerable study may need to be done. Some issues are if this needs to be explicitly done or if the policy can be so written that a common description of the target is also included? Allowing a policy target to retrieve those policies that are relevant to it.

4. Conclusions

Understanding the set of problems facing IP network management in general will be key in defining a comprehensive framework architecture that meets the needs of operators. Additional risks are created by applying new management techniques to the management of IP networks. The consequence of implementing management operations based on architectures that may not be compatible with existing management systems will still need to be explored.

Given that many network devices in IP networks are making routing decisions based on information received via routing protocols it seems sensible that they also make QoS decisions in a similar fashion.

Historically the broader the scope of a network management standardization effort the less likely it has been to succeed. Management standardization efforts must be careful to have clearly defined goals and requirements less they to experience the same fate as previous such efforts.

As IP continues to extend it's concept of service beyond that of best effort to include, among other things, differentiate treatment of packets, it will become increasingly necessary to have mechanisms capable of supporting these extensions. Efforts to define a common management model and framework have proven to be historically elusive. Information models, whether they be traditional or rules-based, must address these past problems. The desire to keep a competitive advantage, and the reality that a common model, to be truly common, will not provide sufficient detail to fully manage a device, has often slowed the acceptance on the part of equipment vendors to this approach.

As IP continues to extend it's concept of service beyond that of best effort to include, among other things, differentiate treatment of packets it will become increasingly necessary to have mechanisms capable of supporting these extensions.

5. Security Considerations

The exchange of management information in a network is one of the most sensitive from a security perspective. Management protocols must address security to insure the integrity of the data. A management architecture must provide for security considerations from its inception to insure the authenticity of the information provider and that the security mechanisms not be so cumbersome as to make them not feasible to implement.

6. Reference

- [1] Michael Eder, Sid Nag, Raj Bansal, "IP Service Management Framework", Work in Progress, October 1999.
- [2] Hugh Mahon, Yoram Bernet, and Shai Herzog, "Requirements for a Policy Management System", Work in Progress.
- [3] Yavatkar, R., Pendarakis, D. and R. Guerin, "A Framework for Policy-based Admission Control", RFC 2753, January 2000.
- [4] Huston, G., "Next Steps for the IP QoS Architecture", RFC 2990, November 2000.
- [5] McCloghrie, K. and M. Rose, "Management Information Base for Network Management of TCP/IP-based internets" RFC 1156, May 1990.

7. Authors' Addresses

Michael Eder
Nokia
5 Wayside Road
Burlington, MA 01803

EMail: michael.eder@nokia.com

Sid Nag
PO Box 104
Holmdel, NJ 07733

EMail: thinker@monmouth.com

8. Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

