

Network Working Group
Request for Comments: 3139
Category: Informational

L. Sanchez
Megisto
K. McCloghrie
Cisco
J. Saperia
JDS Consultant
June 2001

Requirements for Configuration Management of IP-based Networks

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

This memo discusses different approaches to configure networks and identifies a set of configuration management requirements for IP-based networks.

Table of Contents

1.0	Introduction	2
1.1	Motivation, Scope and Goals of this document	2
1.2	Requirements Terminology	3
1.3	Audience	3
1.4	Definition of Technical Terms.	3
2.0	Statement of the Problem	4
3.0	Requirements for an IP-based Configuration Management System .	7
4.0	Security Considerations	9
	Acknowledgments	9
	References.	9
	Authors' Addresses.	10
	Full Copyright Statement.	11

1.0 Introduction

1.1 Motivation, Scope and Goals of this document

A number of IETF working groups have introduced new technologies which offer integrated and differentiated services. To support these new technologies, working group members found that they had new requirements for configuration of these technologies. One of these new requirements was for the provisioning (configuration) of behavior at the network level.

An example of this type of configuration would be instructing all routers in a network to provide 'gold' service to a particular set of customers. Depending on the specific network equipment and definition of 'gold' service, this configuration request might translate to different configuration parameters on different vendors equipment and many individual configuration commands at the router. This higher level of configuration management has come to commonly be known as policy based management.

Working groups associated with these new technologies believed that the existing SNMP based management framework, while adequate for fault, configuration management at the individual instance (e.g., interface) level, performance and other management functions commonly associated with it, was not able to meet these new needs. As a result they began working on new solutions and approaches.

COPS [COPS] for RSVP [RSVP] provides routers with the opportunity to ask their Policy Server for an admit/reject decision for a particular RSVP session. This model allows routers to outsource their resource allocation decisions to some other entity. However, this model does not work with DiffServ [DSARCH] where there is no signalling protocol. Therefore, the policies that affect resource allocation decisions must be provisioned to the routers. It became evident that there was a need for coordinating both RSVP-based and DiffServ-based policies to provide end2end QoS. Working groups began to extend and leverage approaches such as COPS for RSVP to support Diffserv policies. This gave birth to COPS-PR [COPS-PR].

These extensions caused concern that the IETF was about to develop a set of fragmented solutions which were locally optimized for specific technologies and not well integrated in the existing Internet Management Framework. The concern prompted some of the Area Directors associated with the Operations and Management, Transport and General areas, and some IAB members to organize a two day meeting in mid September 1999. The primary purpose of the meeting was to examine the requirements for configuration management and evaluate the COPS/PIB and SNMP/MIB approaches in light of these requirements.

At the end of the two day meeting there was no consensus on several issues and as a result a number of 'design teams' were created. This document is the output of the design team chartered with the identification of a global set of configuration management requirements. This document has benefited from feedback received during the Configuration Management BOF that took place on November 11, 1999 during the 46th IETF in Washington DC, USA. The document has also benefited from comments sent to the confmgt@ops.ietf.org mailing list.

1.2 Requirements Terminology

Keywords "MUST", "MUST NOT", "REQUIRED", "SHOULD", "SHOULD NOT" and "MAY" that appear in this document are to be interpreted as described in RFC 2119 [Bra97].

1.3 Audience

The target audience for this document includes system designers, implementers of network configuration and management technology and others interested in gaining a general background understanding of the issues related to configuration management in general, and in the Internet in particular along with associated requirements. This document assumes that the reader is familiar with the Internet Protocol, related networking technology, and general network management terms and concepts.

1.4 Definition of Terms

Device-Local Configuration

Configuration data that is specific to a particular network device. This is the finest level of granularity for configuring network devices.

Network-Wide Configuration

Configuration data that is not specific to any particular network device and from which multiple device-local configurations can be derived. Network-wide configuration provides a level of abstraction above device-local configurations.

Configuration Data Translator

A function that transforms Configuration Management Data (high-level policies) or Network-wide configuration data (middle-level policies) into device local configurations (low-level policies) based on the

generic capabilities of network devices. This function can be performed either by devices themselves or by some intermediate entity.

2.0 Statement of the Problem

Configuring large networks is becoming an increasingly difficult task. The problem intensifies as networks increase their size, not only in terms of number of devices, but also with a greater variety of devices, with each device having increasing functionality and complexity. That is, networks are getting more complex in multiple dimensions simultaneously (number of devices, time scales for configuration, etc.) making the task of configuring these more complex.

In the past, configuring a network device has been a three step process. The network operator, engineer or entity responsible for the network created a model of the network and its expected behavior. Next, this (model + expected behavior) was formalized and recorded in the form of high-level policies. Finally, these policies were then translated into device-local configurations and provisioned into each network device for enforcement.

Any high-level policy changes (changes in the network topology and/or its expected behavior) needed to be translated and provisioned to all network devices affected by the change. Figure 1 depicts this model and shows how high-level policies for a network could be translated into four device-local configurations. In this model, network operators or engineers functioned as configuration data translators; they translated the high-level policies to device-local configuration data.

A configuration data translator could take the topology independent behavior description such as high-level policies (first input source) combine it with topology information (second input source) as well as status/performance/monitoring information (third input source) to derive device-local configurations. Note that there could be several configuration data translators operating in tandem on a set of devices. However, there could be only one configuration data translator operating at a particular device at any given instance.

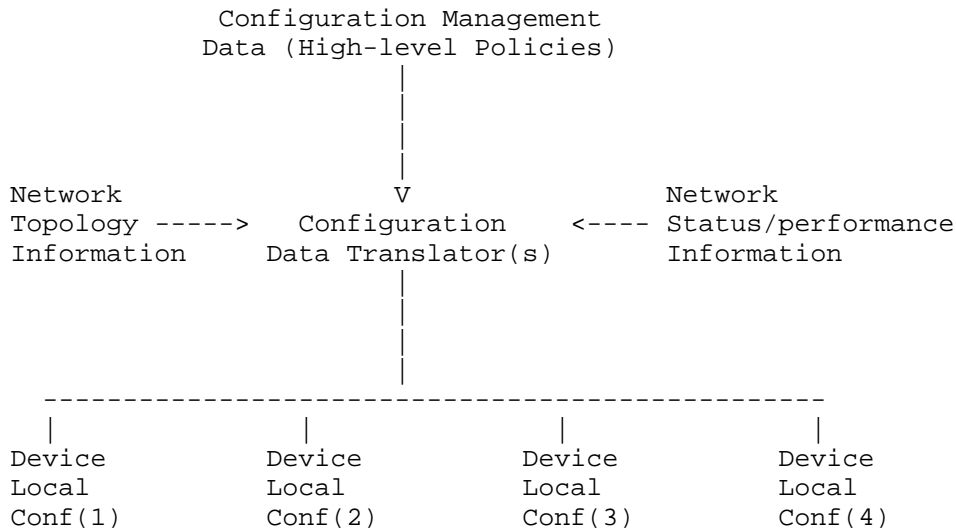


Figure 1. Current model for configuring network devices.

Historically, network operators and engineers used protocols and mechanisms such as SNMP and CLI applications to provision or configure network devices. In their current versions, these mechanisms have proven to be difficult to use because of their low-level of granularity and their device-specific nature. This problem is worse when provisioning multiple network devices requiring large amounts of configuration data.

It is evident that network administrators and existing configuration management software can not keep up with the growth in complexity of networks and that an efficient, integrated configuration management solution is needed. Several IETF Working Groups working on this problem converged into adding a layer of abstraction to the traditional configuration management process described in figure 1. Figure 2 depicts this process after the layer of abstraction is added. As in the previous figure, first the network operator, engineer or entity responsible for the network creates a model of the network and its expected behavior. This is formalized and recorded in the form of high-level policies.

These policies are combined with topology information as well as status/performance information to generate network-wide configuration data. These middle level-policies are simpler to manage and represent behaviors shared by multiple network devices.

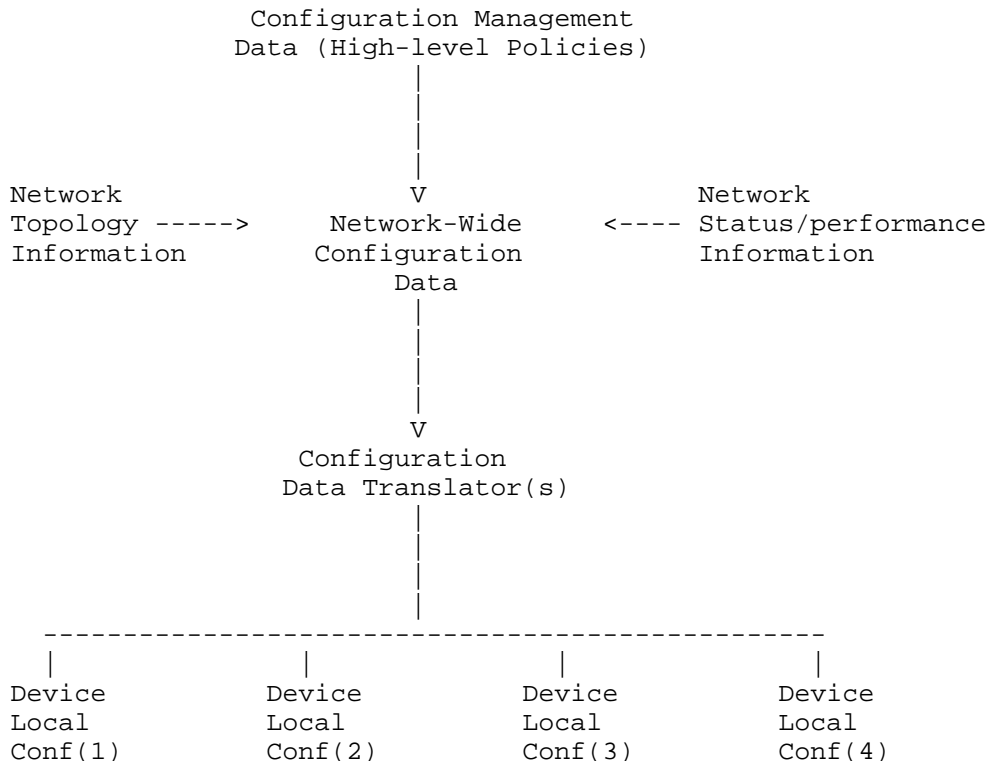


Figure 2. Proposed model for configuring network devices.

Device local configurations are generated by automated configuration data translators and are supplied to each network device for enforcement. Note how this model only describes the function of the configuration data translators and it does not dictate its functional location. This is to say that translators may reside outside of the devices (as it was the case in figure 1 since they were humans) or may be possibly collocated with each device.

As in the previous model, any high-level policy changes (changes in the network topology and/or its expected behavior) needs to be propagated to all network devices affected by the change. However, in the configuration model depicted in figure 2 network operators and engineers can specify the behavior of the network in a simplified manner reducing the amount of device specific knowledge needed.

One should keep in mind that in some cases per instance device local configuration is needed in network devices. An integrated solution MUST allow room for this. Also, the introduction of automated configuration data translators assumes that all information needed to

make an error free conversion of network-wide configuration data into device-local configuration data is available. In the event that such data is not available the solution MUST detect this and act accordingly.

3.0 Requirements for an IP-based Configuration Management System

All IETF WGs active in this area agrees upon the following requirements for configuration management. An integrated configuration management solution MUST:

- 1) provide means by which the behavior of the network can be specified at a level of abstraction (network-wide configuration) higher than a set of configuration information specific to individual devices,
- 2) be capable of translating network-wide configurations into device-local configuration. The identification of the relevant subset of the network-wide policies to be down-loaded is according to the capabilities of each device,
- 3) be able to interpret device-local configuration, status and monitoring information within the context of network-wide configurations,
- 4) be capable of provisioning (e.g., adding, modifying, deleting, dumping, restoring) complete or partial configuration data to network devices simultaneously or in a synchronized fashion as necessary,
 - 4a) be able to provision multiple device-local configurations to support fast switch-overs without the need to down-load potentially large configuration changes to many devices,
- 5) provide means by which network devices can send feedback information (configuration data confirmation, network status and monitoring information, specific events, etc.) to the management system,
- 6) be capable of provisioning complete or partial configuration data to network devices dynamically as a result of network specific or network-wide events,
- 7) provide efficient and reliable means compared to current versions of today's mechanisms (CLI, SNMP) to provision large amounts of configuration data,

- 8) provide secure means to provision configuration data. The system must provide support for access control, authentication, integrity-checking, replay- protection and/or privacy security services. The minimum level of granularity for access control and authentication is host based. The system SHOULD support user/role based access control and authentication for users in different roles with different access privileges,
- 9) provide expiration time and effective time capabilities to configuration data. It is required that some configuration data items be set to expire, and other items be set to never expire,
- 10) provide error detection (including data-specific errors) and failure recovery mechanisms (including prevention of inappropriately partial configurations when needed) for the provisioning of configuration data,
- 11) eliminate the potential for mis-configuration occurring through concurrent shared write access to the device's configuration data,
- 12) provide facilities (with host and user-based authentication granularity) to help in tracing back configuration changes,
- 13) allow for the use of redundant components, both network elements and configuration application platforms, and for the configuration of redundant network elements.
- 14) be flexible and extensible to accommodate future needs. Configuration management data models are not fixed for all time and are subject to evolution like any other management data model. It is therefore necessary to anticipate that changes will be needed, but it is not possible to anticipate what those changes might be. Such changes could be to the configuration data model, supporting message types, data types, etc., and to provide mechanisms that can deal with these changes effectively without causing inter-operability problems or having to replace/update large amounts of fielded networking devices,
- 15) leverage knowledge of the existing SNMP management infrastructure. The system MUST leverage knowledge of and experience with MIBs and SMI.

Security Considerations

This document reflects the current requirements that the IETF believes configuration management systems MUST have to properly support IP-based networks. The authors believe that a configuration management system MUST provide mechanisms by which one can ascertain the integrity and authenticity of the configuration data at all times. In some cases the privacy of the data is important therefore configuration management system MUST provide facilities to support this services as required not only while the data is stored but also during provisioning or reception. Requirements eight and twelve capture the required security services.

Acknowledgments

The authors thank Juergen Schoenwaelder for his contributions to this document. The authors also thank Walter Weiss and Andrew Smith for providing feedback to early versions of this document. Finally, the authors thank the IESG for motivating and supporting this work.

References

- [Bra97] Bradner, S., "Key Words for use in RFCs to indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [COPS] Boyle, J., Cohen, R., Durham, D., Herzog, S., Rajan, R. and A. Sastry, "The COPS (Common Open Policy Service) Protocol", RFC 2748, August 1999.
- [RSVP] Braden, R., Editor, et al., "Resource ReSerVation Protocol (RSVP) Version 1 - Functional Specification", RFC 2205, September 1997.
- [COPS-RSVP] Boyle, J., Cohen, R., Durham, D., Herzog, S., Rajan, R. and A. Sastry, "COPS usage for RSVP", RFC 2749, June 1999.
- [COPS-PROV] Chan, K., Seligson, J., Durham, D., Gai, S., McCloghrie, K., Herzog, S., Reichmeyer, F., Yavatkar, R. and A. Smith, "COPS Usage for Policy Provisioning (COPS-PR)", RFC 3084, March 2001.

Authors' Addresses

Keith McCloghrie
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA

Phone: +1 (408) 526-5260
EMail: kzm@cisco.com

Luis A. Sanchez
Megisto Systems
20251 Century Boulevard
Germantown, MD 02138
USA

Phone: +1 (301) 444-1747
EMail: lsanchez@megisto.com

Jon Saperia
JDS Consulting, Inc.
174 Chapman Street
Watertown, MA 02472
USA

Phone: +1 (617) 744-1079
EMail: saperia@jdscons.com

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.