Distributed Management Framework

<draft-ietf-disman-framework-02.txt>

August 23, 1998

Authors:

Andy Bierman
Cisco Systems
abierman@cisco.com

Maria Greene
Ascom Nexion
greene@nexen.com

Bob Stewart
Cisco Systems
bstewart@cisco.com

Steve Waldbusser
International Network Services (INS)
waldbusser@ins.com

1. Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

To view the entire list of current Internet-Drafts, please check the "1id-abstracts.txt" listing contained in the

Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), ftp.nordu.net (Northern Europe), ftp.nis.garr.it (Southern Europe), munnari.oz.au (Pacific Rim), ftp.ietf.org (US East Coast), or ftp.isi.edu (US West Coast).

2. Abstract

This memo defines a distributed management architecture for use with the SNMP network management architecture.

This memo does not specify a standard for the Internet community.

3. Overview

Distributed Management is the delegation of control from one management station to another. While the SNMP Network Management Framework does not specifically address distributed management, this function is desired and has been implemented and deployed in the internet using proprietary architectures. It is desired that there be a standard upon which to promote interoperability, as well as a common framework upon which various systems can be built.

The goals of distributed management are:

- Scalability through Distribution
  In order to build network management systems that have the power to manage very large networks, it is important to reduce bottlenecks in the management system. Therefore, a distributed systems approach is often helpful when building large management systems. A distributed approach is often very effective at reducing load on the central management station, and may be effective at reducing the load that the central management station places on backbone networks. However, a distributed approach usually has no benefit in reducing load on remote networks and has no benefit in reducing load on management agents. Further, in a distributed data collection architecture, if all data collected is eventually forwarded to the central management station (without aggregation or compression), then no benefit in backbone load or central management station load should be expected (except perhaps to time-shift this load to a time of excess capacity, at the expense of a lack of timeliness of data.

- Disconnected or Low-Bandwidth Operation
  There are sometimes conditions when a management station will not be in constant contact with all portions of the managed network. This is sometimes by design in an attempt to lower communications costs (especially when communicating over a WAN or dialup link), or by accident as network failures affect the communications between the management station and portions of the managed network.

  For this reason, a distributed management station will be configured to perform network management functions, even when communication with the management station may not be

possible or efficient. The distributed management station may then attempt to notify the management station when an exceptional condition occurs. Thus, even in circumstances where communication with the distributed management station is not continuous, network management operations may run continuously, communicating with the management station conveniently and efficiently, on an as-needed basis.

- o Mirroring organization boundaries and processes
  Business organizations are typically set up in a hierarchical fashion. Multiple people in the hierarchy have different roles, responsibilites, and authority. The network management system often has to be configured to match this organization. Distributed network managers can be set up in a hierarchy that matches the roles of various people in the organization.

- o Promotes a modular system architecture
  A modular system architecture allows flexibility and re-usability of network management components. This also enables a multi-vendor approach to building management systems. A distributed network management system with well-defined interfaces creates this modular scheme.

  This document defines an architectural framework for standards-based distributed management

4. Relationship to the SNMP Management Framework

A distributed network management station is a management station that receives requests from another manager and executes those requests by performing management operations on agents or other managers. Note that these requests may take a long time to execute, and may be registered indefinitely. This framework uses the services of the SNMP Management Framework.

4.1. The SNMP Management Framework

The SNMP Management Framework presently consists of five major components:

- o An overall architecture, described in RFC 2271 [1].

- Mechanisms for describing and naming objects and events for the purpose of management. The first version of this Structure of Management Information (SMI) is called SMIv1 and described in RFC 1155 [2], RFC 1212 [3] and RFC 1215 [4]. The second version, called SMIv2, is described in RFC 1902 [5], RFC 1903 [6] and RFC 1904 [7].

- Message protocols for transferring management information. The first version of the SNMP message protocol is called SNMPv1 and described in RFC 1157 [8]. A second version of the SNMP message protocol, which is not an Internet standards track protocol, is called SNMPv2c and described in RFC 1901 [9] and RFC 1906 [10]. The third version of the message protocol is called SNMPv3 and described in RFC 1906 [10], RFC 2272 [11] and RFC 2274 [12].

- Protocol operations for accessing management information. The first set of protocol operations and associated PDU formats is described in RFC 1157 [8]. A second set of protocol operations and associated PDU formats is described in RFC 1905 [13].

- A set of fundamental applications described in RFC 2273 [14] and the view-based access control mechanism described in RFC 2275 [15]. Managed objects are accessed via a virtual information store, termed the Management Information Base or MIB. Objects in the MIB are defined using the mechanisms defined in the SMI. This memo specifies a MIB module that is compliant to the SMIv2. A MIB conforming to the SMIv1 can be produced through the appropriate translations. The resulting translated MIB must be semantically equivalent, except where objects or events are omitted because no translation is possible (use of Counter64). Some machine readable information in SMIv2 will be converted into textual descriptions in SMIv1 during the translation process. However, this loss of machine readable information is not considered to change the semantics of the MIB.

5. Distributed Management Framework

The distributed management framework consists of applications and services.

A distributed management application performs some management function, often by monitoring and controlling managed elements. These applications perform functions such as threshold polling, historical data polling, or discovery. The specifications and communications protocols of some of these applications will be defined by standards, while others will be enterprise specific.

Distributed management services are a collection of services that make up the run-time environment for the distributed management application. These services are crucial to ensuring the usability, scalability, and efficiency of the distributed management applications that depend on them. Some of these services are mandatory, while others are optional. Further, even the mandatory services allow varying depths of support, as described further, below.

Each of these services is made available through a MIB interface.

The purpose of these services is to provide true enterprise management that allows distributed management functions to be used on an enterprise-wide basis without undue amounts of administrative difficulty. These services also make a set of applications more efficient because the service can perform functions or store information once for all applications on the local system. Further, less work is required to be put into writing the application because some of the core functions are performed elsewhere.

5.1. Known Systems

The Known Systems service provides a list of all systems which the distributed management system knows about and is prepared to perform management operations upon. This list may be populated by a continuously-running auto-discovery process, it may be populated with SNMP SET requests, or it may be a static list dictated by the system.

This service also stores type and attribute information for each of the known systems.

The purpose of this service is to allow a management application to be executed on a list of systems so that true enterprise management is possible rather than more ad-hoc management functions. Further, the targets of management applications can be configured separately from the applications to reduce administrative burden as the list of known systems changes.

An example of a known systems list is a list of all systems at a site discovered by the autodiscovery module, along with various addressing, type, and attribute information for each.

5.2. Management Domains

The Management Domains service provides a list of known systems which is a subset of the entire list of known systems. The subset is defined by a rule, or filter, which selects only the known systems that match or those that don't match certain criteria.

These domains are multiple, potentially overlapping, sets of devices based on (human) organizational boundaries that define the extents over which management operations should be performed.

The purpose of this service is to allow a management application to be executed on a list of known systems within a particular domain. Further, as the domains change over time, the application will automatically be kept up to date and will only monitor the correct systems.

An example of a management domain is the subset of all known systems that is on the engineering LAN.

5.3. Management Operations Targets

The Management Operations Targets service provides a list of known systems in a set of domains that match certain criteria that indicate that it makes sense to perform a particular management function on them.

The purpose of this service is to direct management operations to be performance only on those systems where that operation would make sense. Because this is described as a filter, there are no static configuration requirements that would be unacceptable in a dynamically changing network environment.

An example of a management operation target list is the subset of all known routers on the engineering LAN.

5.4. Credential Delegation

The Credential Delegation Service allows credentials of a network management user to be delegated to a distributed management application so that it can perform secure operations on behalf of that user.

Fundamental to this distributed management architecture is the notion that distributed management operations must not run with the credentials of the distributed manager. To do so would require that the authorization of these credentials (or subsets of this authorization) would need to be apportioned to users of that distributed manager in a pre-defined and secure way. This would require the creation of a access control architecture mirroring the SNMP View-Based Access Control architecture that would control what subsets of authority are available to what users. The existing View-Based Access Control mechanism was not designed for this task and is not appropriate. Further, it would require that the distributed manager be configured in a way that was consistent with the access control policy embodied in the managed systems. This would be extremely problematic because:

1. This would require a massive amount of configuration to be replicated on the distributed manager

2. If any part of this configuration on the distributed manager is inconsistent with that on the managed systems, a security hole could be exposed.

Because it is assumed that the distributed manager adds no credentials to management operations, when a manager wishes to configure a distributed manager to perform secure transactions on its behalf, it must download to the distributed manager the appropriate credentials to be placed in secure SNMP messages, identifying them as the manager. A credential contains at

least the securityModel, securityName, securityLevel, authentication and privacy keys, and an indication of which management targets the credential should be used for.

5.4.1. Definitions

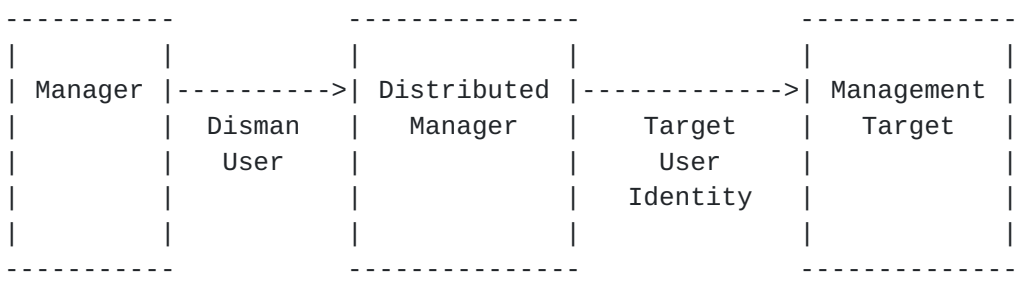

1. Disman User - The user whose credentials are used to configure the distributed manager for an operation and to download credentials for that operation. There is no relationship implied between the disman user and the user(s) who's credentials are installed (in other words, "joe" can install credentials for "ops-center-east" as well as "joe").

2. Target User Identity - The user identity used in SNMP messages between the distributed manager and management targets.

3. Credential - The set of secrets that are transferred to the distributed manager giving it the authority to act as an identity.

4. Owner - The disman user who sets up a distributed management function, including the credentials for the function.

5. Invoker - The user who invokes a previously setup distributed management function. The owner may choose to allow others to invoke a function, potentially allowing that function to operate with the owner's credentials (of course the owner would want to tightly constrain what the function was configured to perform).

6. Invokation Identity - The identity of the credentials a function is operating with. These may be of the owner, of the invoker, or possibly the union of both

credentials.

Because multiple Disman Users will have access to a Distributed Manager, the Credential Delegation Service will be responsible for ensuring that credentials are only used by authorized users. The Credential Delegation Service will include:

1. Credential Store - a MIB in which to transfer and store credentials

2. MIB prototype - a prototype MIB fragment that will be added to disman functions that wish to use the Credential Store

3. Access Control Policy - a policy for configuration of the VACM MIB for use with the Credential Delegation Service. This will limit access to the credential store.

5.5. Delegation Control

The Delegation Control Service provides functions that limit the resource usage of a distributed management application that has had control delegated to it.

Network management applications are often responsible for adding stress on the network and causing problems. Examples are excessive polling load on slow-speed networks or on router CPUs. This problem will become even more dangerous when network management functions are delegated to distributed management stations.

Policies need to be configured that limit average and burst polling, notification, and broadcast rates. These rates may be defined for the sending system as a whole, per end node, or per management application on the sending system.

It is also important to have a "Deadman's switch" so that delegated applications will not continue indefinitely if their "sponsor" has forgotten about them.

5.6. Scheduling

The Scheduling Service allows the execution of distributed management applications to be controlled so that they run at a particular time, periodically, or based on the occurance of another event.

5.7. Reliable Notification

The Reliable Notification Service provides services that ensure that notifications are received correctly.

For example, all the information that describes an event may not fit within a single PDU, so an eventID varbind is defined that associates multiple PDU's as describing the same event. It is also necessary to know if the entire notification has been received or if more PDU's are still outstanding.

Further, a receiving management station must be given more information so that it can distinguish duplicated inform PDU's because events are not idempotent. No rule makes it mandatory for the requestID to be unique for all PDUs sent from a system.

In addition, a logging mechanism provides for retrieval of notifications after a communications failure.

5.8. Notification Destinations

The Notification Destination Service provides services for configuring destinations for notifications.

When management functions are delegated and MLMs are given the autonomy to generate notifications, they need to be configured as to where the notifications should be sent. Additionally, retry counts and numbers need to be configured. Average and burst notification rates need to be defined.

6. Acknowledgments

This document was produced by the IETF Distributed Network Management Working Group.

7. References

[1] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing SNMP Management Frameworks", RFC 2271, Cabletron Systems, Inc., BMC Software, Inc., IBM T. J. Watson Research, January 1998

[2] Rose, M., and K. McCloghrie, "Structure and Identification of Management Information for TCP/IP-based Internets", RFC 1155, Performance Systems International, Hughes LAN Systems, May 1990

[3] Rose, M., and K. McCloghrie, "Concise MIB Definitions", RFC 1212, Performance Systems International, Hughes LAN Systems, March 1991

[4] M. Rose, "A Convention for Defining Traps for use with the SNMP", RFC 1215, Performance Systems International, March 1991

[5] Case, J., McCloghrie, K., Rose, M., and S. Waldbusser, "Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMPv2)", RFC 1902, SNMP Research,Inc., Cisco Systems, Inc., Dover Beach Consulting, Inc., International Network Services, January 1996.

[6] Case, J., McCloghrie, K., Rose, M., and S. Waldbusser, "Textual Conventions for Version 2 of the Simple Network Management Protocol (SNMPv2)", RFC 1903, SNMP Research, Inc., Cisco Systems, Inc., Dover Beach Consulting, Inc., International Network Services, January 1996.

[7] Case, J., McCloghrie, K., Rose, M., and S. Waldbusser, "Conformance Statements for Version 2 of the Simple Network Management Protocol (SNMPv2)", RFC 1904, SNMP Research, Inc., Cisco Systems, Inc., Dover Beach Consulting, Inc., International Network Services, January 1996.

[8] Case, J., Fedor, M., Schoffstall, M., and J. Davin, "Simple Network Management Protocol", RFC 1157, SNMP Research, Performance Systems International, Performance Systems International, MIT Laboratory for Computer Science, May 1990.

[9] Case, J., McCloghrie, K., Rose, M., and S. Waldbusser, "Introduction to Community-based SNMPv2", RFC 1901, SNMP Research, Inc., Cisco Systems, Inc., Dover Beach Consulting, Inc., International Network Services, January 1996.

[10] Case, J., McCloghrie, K., Rose, M., and S. Waldbusser, "Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2)", RFC 1906, SNMP Research, Inc., Cisco Systems, Inc., Dover Beach Consulting, Inc., International Network Services, January 1996.

[11] Case, J., Harrington D., Presuhn R., and B. Wijnen, "Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)", RFC 2272, SNMP Research, Inc., Cabletron Systems, Inc., BMC Software, Inc., IBM T. J. Watson Research, January 1998.

[12] Blumenthal, U., and B. Wijnen, "User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)", RFC 2274, IBM T. J. Watson Research, January 1998.

[13] Case, J., McCloghrie, K., Rose, M., and S. Waldbusser, "Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)", RFC 1905, SNMP Research, Inc., Cisco Systems, Inc., Dover Beach Consulting, Inc., International Network Services, January 1996.

[14] Levi, D., Meyer, P., and B. Stewart, "SNMPv3 Applications", RFC 2273, SNMP Research, Inc., Secure Computing Corporation, Cisco Systems, January 1998

[15] Wijnen, B., Presuhn, R., and K. McCloghrie, "View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)", RFC 2275, IBM T. J. Watson Research, BMC Software, Inc., Cisco Systems, Inc., January 1998

Table of Contents