



IBM Tivoli Training

Netcool/Proviso 4.4.3.1

High availability manager (HAM) concepts



© 2009 IBM Corporation
Updated January 26, 2009

IBM Tivoli Netcool®/Proviso® 4.4.3.1 is a premier network performance management tool that provides quick, comprehensive, reporting of network performance.

Netcool/Proviso collects network performance data from a number of different sources including Simple Network Management Protocol (SNMP).

Netcool/Proviso 4.4.3.1 provides an optional process for SNMP collection that ensures high availability in the face of SNMP collector host outages.

Objectives

- Upon completion of this module, you should be able to:
- Describe the concept of High Availability Manager (HAM) for SNMP Collectors
- Describe the use of HAM
- Describe the components of HAM
- Recognize a simple and complex HAM cluster

Upon completion of this module, you should be able to:

Describe the concept of High Availability Manager (HAM) for SNMP Collectors

Describe the use of HAM

Describe the components of HAM

Recognize a simple and complex HAM cluster

Assumptions

Before implementing HAM:

- You must have Netcool/Proviso 4.4.3.1 installed and running
- You have root access in order to use Netcool/Proviso's Deployer
- You know how to use Netcool/Proviso's Topology Editor

Before implementing HAM:

You must have Netcool/Proviso 4.4.3.1 installed and running

You must have root access in order to use Netcool/Proviso Deployer

You must know how to use Netcool/Proviso's Topology Editor

HAM overview

- HAM is an optional tool available with the Netcool/Proviso 4.4.3.1 release.
- HAM can be used to provide redundant SNMP collection paths.
- HAM offers options for organizing failover clusters.
- HAM is configured using the Netcool/Proviso topology editor.

HAM is an optional tool that is included in the Netcool/Proviso 4.4.3.1 release. It can be used to provide redundant SNMP collection paths in the event of a collector process outage.

HAM offers several options for organizing failover clusters and is configured using the Netcool/Proviso topology editor.

Definitions

- **Collector profile** is a set of properties that identify the collector; the collector number, the polling interval, and the output directory
- **Collector process** is the DataLoad collection component running on a host
- **Collector host** is the host where the collector process is running
- **Primary host** has collection responsibilities
- **Fixed spare** is a host without collection responsibilities
- **Floating spare** is a host that has primary collection responsibilities but can assume the role of a spare
- **Cluster** is a set of hosts that act as collection primaries and spares for one or more collector profiles

There are a number of terms used in conjunction with HAM. They are defined as follows:

The **Collector profile** is a set of properties that identify the collector. These include the collector number, the polling interval, and the output directory.

A **Collector process** is the DataLoad collection component running on a host.

A **Collector host** is the host where the collector process is running.

A **Primary host** has collection responsibilities.

A **Fixed spare** is a host without collection responsibilities.

A **Floating spare** is a host that has primary collection responsibilities but can assume the role of a spare. It becomes a spare after the host has recovered from an outage and is no longer bound to a collector profile.

A **Cluster** is a set of hosts that act as collection primaries and spares for one or more collector profiles.

HAM basics

- HAM is a separate process running on a host within the Netcool/Proviso installation.
- Normally a collector profile and collector process are inseparably bound.
- Without HAM if a collector process fails the collector itself will not be available until the collector process returns to service.
- With HAM a collector profile can be bound to a specified spare when the primary collector process fails.
- HAM organizes primary collector processes and spares in clusters.
- Each primary collector process has a defined set of spares and is known as a resource pool.

There are some basic HAM concepts that need to be explained before examining HAM implementations. HAM is a specific process that runs on a host within the Netcool/Proviso installation. It is separate from the collection processes.

Without HAM a collector profile and collector process is inseparably bound. In that case if a collector process fails the collector itself will not be available until the collector process returns to service.

With HAM, a collector profile and collector process is not inseparably bound. If a collector process fails, the collector profile can be bound to a specified spare.

HAM organizes primary collector processes and spares in clusters. Each primary collector process has a defined set of spares.

Simple HAM implementation

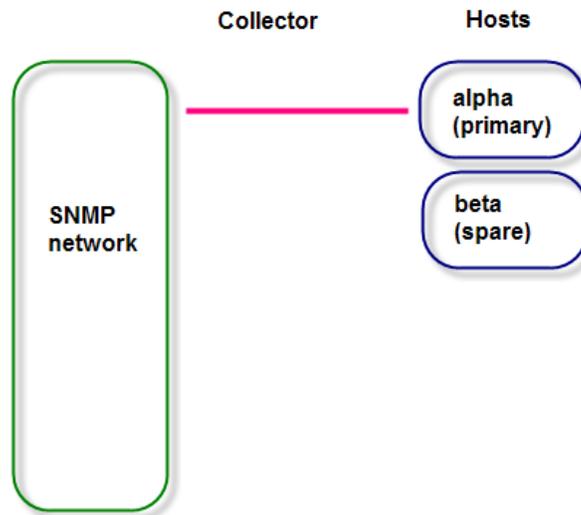
- HAM can be deployed on a single collector profile with a single primary collector process and one fixed spare.
- The HAM unbinds the collector profile from the primary collector process and binds it to the designated spare.
- When the primary collector process recovers it must be manually reassigned as primary to the collector profile to maintain failover coverage.

In a simple HAM implementation a cluster is created for a collector profile and includes a primary collector process and one fixed spare.

Upon failure of the primary collector process, HAM unbinds the collector profile from the primary and binds it to the designated spare.

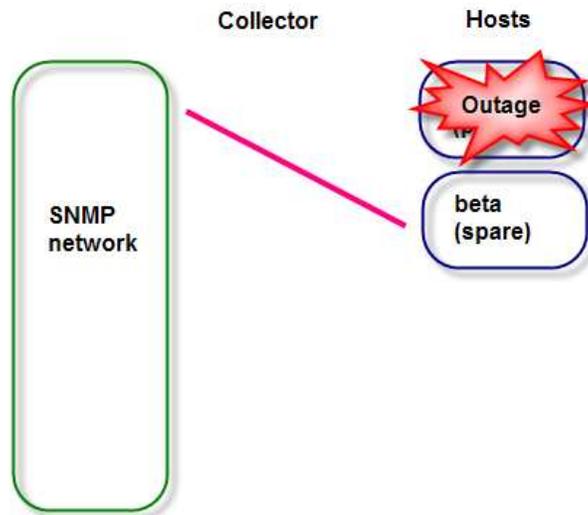
When the primary collector process recovers it must be manually reassigned as primary to the collector profile to maintain failover coverage.

HAM – Primary and one dedicated spare



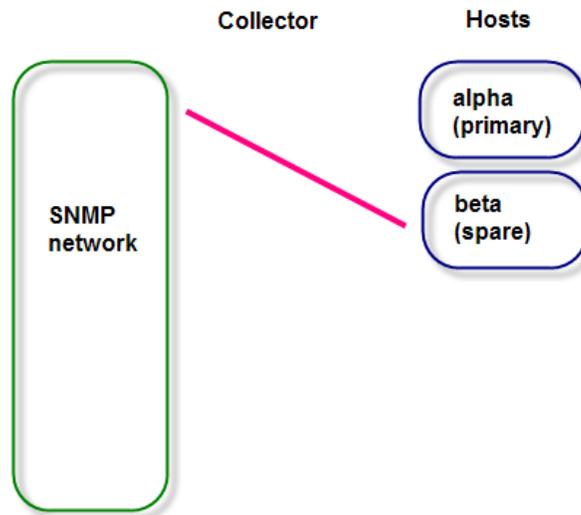
In this example of a simple HAM implementation, the collector profile is bound to the collector process running on the primary host alpha. The fixed spare is the host beta. This shows the normal state of the collector profile and collector process.

HAM – Outage on primary



The primary collector process is offline because of an outage on host alpha. HAM will unbind the collector profile from the offline primary collector process. It will start the collector process on the fixed spare beta and then will bind the collector profile to beta.

HAM – Outage over



Even after the primary collector process host returns to service, HAM still has the collector profile bound to the fixed spare on host beta.

If the collector process host beta were to experience an outage at this point the collector profile would not be protected.

The Netcool/Proviso administrator must manually unbind the fixed spare from the collector profile. The collector profile then must be rebound to the primary collector process host.

Complex HAM implementation

- There are multiple collector profiles in this cluster.
- Each collector profile has a primary collector process in its resource pool.
- Each primary collector process also serves as a floating spare.
- Each resource pool has a set of floating and fixed spares defined to it.
- Each resource pool has a hierarchy of spare utilization.
- When a primary collector process that is also a floating spare is unavailable it will become a spare when it returns to service.

In the following representation of a complex implementation of HAM:

There are multiple collector profiles in the cluster.

Each collector profile has a primary collector process in its resource pool.

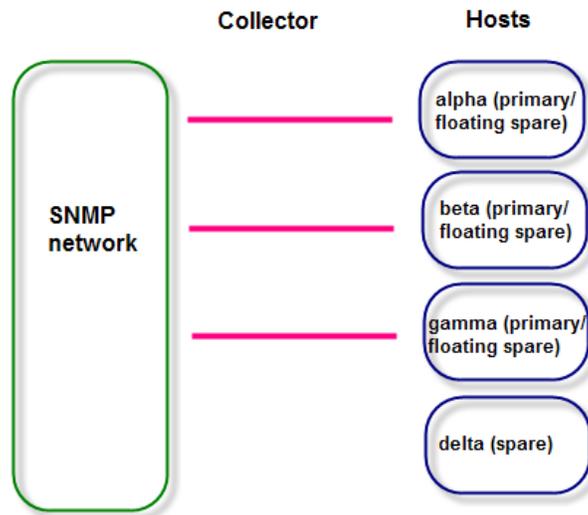
Each primary collector process also serves as a floating spare.

Each resource pool has a set of floating and fixed spares defined to it.

Each resource pool has a hierarchy of spare utilization.

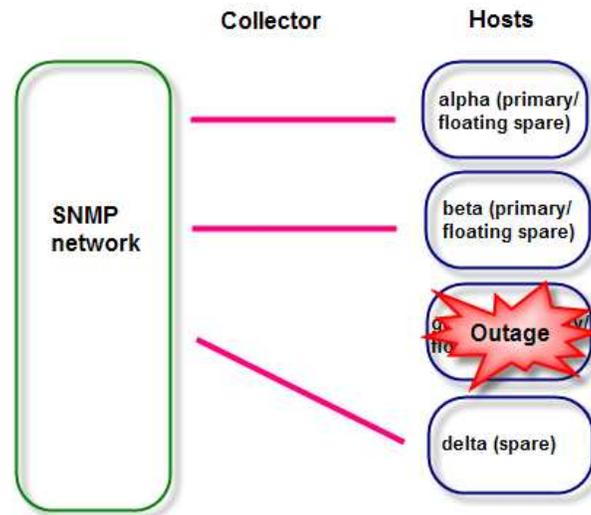
When a primary collector process that is also a floating spare is unavailable it will become a spare when it returns to service.

HAM – Floating spares and dedicated spare



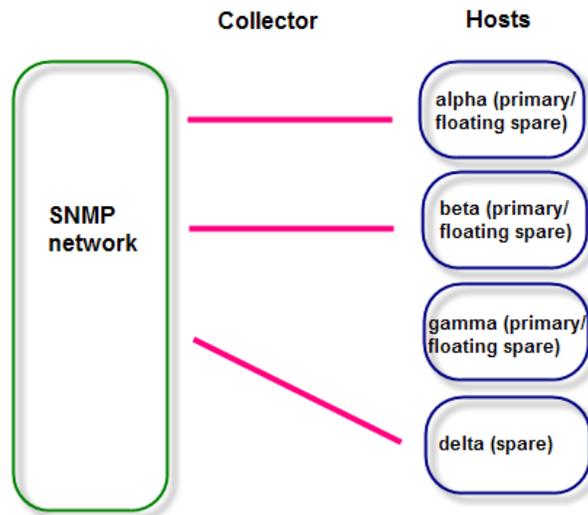
In this cluster there are three collector profiles managed by HAM. The profiles are currently bound to three primary collector process hosts, alpha, beta, and gamma. The three primary collector process hosts are also floating spares. There is a fixed spare host, delta.

HAM – Outage on primary



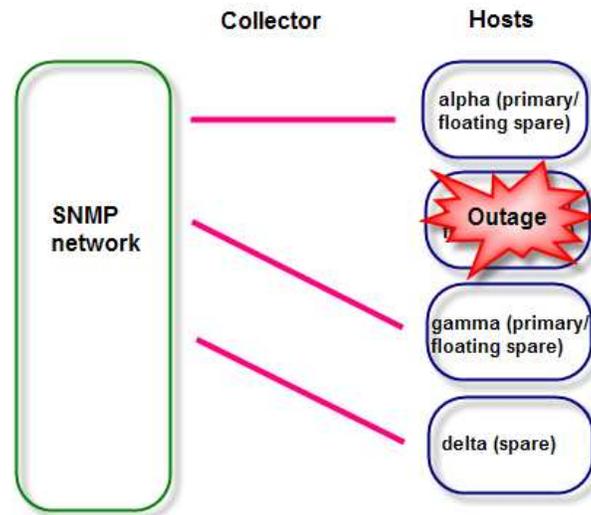
In this view of the cluster the host gamma has experienced an outage. HAM has bound the collector profile from gamma to the fixed spare delta based on how the resource pool was set up.

HAM – Outage over



When the host gamma returns to service it will become a floating spare since the collector profile that was using gamma is now bound to delta.

HAM – Second outage - floating spare



In this view the collector process host beta is in an outage. HAM will now bind the collector profile to the floating spare gamma.

Training roadmap for *Netcool/Proviso*

http://www.ibm.com/software/tivoli/education/edu_prd.html



Copy and paste the link provided into the browser of your choice to explore the training roadmap for Netcool/Proviso.

Summary

- You should now be able to:
- Describe the concept of High Availability Manager (HAM) for SNMP Collectors
- Describe the use of HAM
- Describe the components of HAM
- Recognize a simple and complex HAM cluster

You should now be able to:

Describe the concept of High Availability Manager (HAM) for SNMP Collectors

Describe the use of HAM

Describe the components of HAM

Recognize a simple and complex HAM cluster

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_ham_concepts.ppt

This module is also available in PDF format at: ../ham_concepts.pdf



You can help improve the quality of IBM Education Assistant content by providing feedback.

Trademarks, copyrights, and disclaimers

IBM, the IBM logo, ibm.com, and the following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

Netcool Proviso Tivoli

If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of other IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>

Other company, product, or service names may be trademarks or service marks of others.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2009. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.

