VMware Smart Assurance Topology Split Manager User Guide

VMware Smart Assurance 10.0.0

vmware®

You can find the most up-to-date technical documentation on the VMware website at:

https://docs.vmware.com/

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2019 VMware, Inc. All rights reserved. Copyright and trademark information.

Contents

1 Preface 5

Purpose 5

Audience 5

2 VMware Smart Assurance IP Manager installation directory 6

3 VMware Smart Assurance IP Manager products 7

4 Related documentation 8

Help for command line programs 8

VMware Smart Assurance IP Manager documentation 8

VMware Smart Assurance documentation 9

5 Conventions used in this document 10

Typographical conventions 10

6 Where to get help 12

7 Your comments 13

8 Introduction 14

Overview 14

Topology representation and root-cause analysis 18

Edge connections and proxy devices 18

Architecture 20

Installation 22

Service command 27

Configuration 28

Discovery 29

Invocation 29

Operating modes 29

Full and incremental topology splits 30

Core topology-split functions 30

Trial learning mode 32

Automatic topology synchronization 37

Edge management 39

TSM hook script 42

9 Configuring the Topology Split Manager 45

Overview 45
Configuration directory 47
User configuration parameters 48
User configuration synchronization across installations 49
Methods for modifying user configuration parameters 50
Command sm_tpmgr and example 56
Description of tsm.conf 56
Description of topo-split.conf 61
Description of override-polling-system.txt 62
Tuning Topology Split Manager discovery 62
Specifying domain names for topology split 66
Specifying a list of devices for one-minute polling 66
Enabling topology split 67
Enabling trial learning mode 68
Enabling automatic topology synchronization 69
Enabling edge management 70
Enabling the TSM hook script 71
Configuring security 73

10 Initiating Discovery and Topology Split 74

Overview 74
Topology Split Manager best practices 75
Preparing for discovery 76
Starting the Broker and the Manager s 78
Performing an initial full discovery in trial learning mode 78
Performing an initial full discovery in normal mode 79
Performing an incremental discovery 81
Initiating a discovery 83
Initiating a full topology split 85
Initiating an incremental topology split 86
Using an initial topology split from trial learning mode 87

11 Topology-Split Data Model 90

Data model 90
Topology-split algorithm and data model 92
Topology Split Manager events 93
Topology-split objects and their attributes 94

12 The sm_edit Utility 98

sm_edit 98
sm_edit example 99

Preface

1

As part of an effort to improve its product lines, VMware Smart Assurance periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your VMware technical support professional if a product does not function properly or does not function as described in this document.

Note This document was accurate at publication time. Go to VMware Smart Assurance Online Support docs.vmware.com to ensure that you are using the latest version of this document.

This chapter includes the following topics:

- Purpose
- Audience

Purpose

This document provides description, configuration, and operation information for the VMware Smart Assurance Topology Split Manager, which is a tool that provides a scalable topology-split functionality for very large IP topology domains.

Audience

This document is part of the VMware Smart Assurance IP Manager documentation set. It is intended for:

- System administrators who are configuring and using the Topology Split Manager.
- IT managers who are seeking to better understand the value of the Topology Split Manager.

VMware Smart Assurance IP Manager installation directory

2

In this document, the term BASEDIR represents the location where VMware Smart Assurance software is installed:

- For UNIX, this location is /opt/InCharge/<product>.

  On UNIX operating systems, VMware Smart Assurance IP Manager is, by default, installed to /opt/InCharge/IP/smarts. This location is referred to as BASEDIR/smarts.

  Optionally, you can specify the root of BASEDIR to be something other than /opt/InCharge (on UNIX) but you cannot change the <product> location under the root directory.

  The VMware Smart Assurance System Administration Guide provides detailed information about the directory structure for VMware Smart Assurance software.

VMware Smart Assurance IP Manager products

3

The VMware Smart Assurance IP Manager includes the following products:

- VMware Smart Assurance IP Availability Manager
- VMware Smart Assurance IP Performance Manager
- VMware Smart Assurance IP Availability Manager Extension for NAS

Related documentation

4

In addition to this document, VMware, Inc. provides a Help system for command line programs as well as product documentation.

This chapter includes the following topics:

- Help for command line programs
- VMware Smart Assurance IP Manager documentation
- VMware Smart Assurance documentation

Help for command line programs

Descriptions of command line programs are available as HTML pages. The index.html file, which provides an index to the various commands, is located in the BASEDIR/smarts/doc/html/usage directory.

VMware Smart Assurance IP Manager documentation

The following documents are relevant to users of the VMware Smart Assurance IP Manager:

- *VMware Smart Assurance IP Manager Release Notes*
- *VMware Smart Assurance Installation Guide for SAM, IP, ESM, MPLS, and NPM Managers*
- *VMware Smart Assurance Open Source License and Copyright Information*
- VMware Smart Assurance Failover System User Guide
- *VMware Smart Assurance IP Manager Deployment Guide*
- *VMware Smart Assurance IP Manager Concepts Guide*
- *VMware Smart Assurance IP Manager User Guide*
- *VMware Smart Assurance IP Manager Reference Guide*
- *VMware Smart Assurance IP Manager Troubleshooting Guide*
- *VMware Smart Assurance IP Manager Documentation Portfolio*
- *VMware Smart Assurance IP Manager Certification Matrix*

VMware Smart Assurance documentation

Readers of this guide might find the following related documentation helpful. These documents are updated periodically. Electronic versions of the updated manuals are available on VMware Online Support:

- VMware Smart Assurance System Administration Guide.
- VMware Smart Assurance Common Information Model (ICIM) Reference for Service Assurance Manager.
- VMware Smart Assurance Common Information Model (ICIM) 1.11 Reference for Non-Service Assurance Manager Products.
- VMware Smart Assurance ASL Reference Guide.
- VMware Smart Assurance Perl Reference Guide.
- VMware Smart Assurance MODEL Reference Guide.
- *Smarts Assurance Foundation EMC Data Access API (EDAA) Programmer Guide*.
- *Smart Assurance Foundation Documentation Portfolio*.

Conventions used in this document

5

VMware Smart Assurance uses the following conventions for special notices:

DANGER

DANGER indicates a hazardous situation which, if not avoided, will result in death or serious injury.

WARNING

WARNING indicates a hazardous situation which, if not avoided, could result in death or serious injury.

CAUTION

CAUTION used with the safety alert symbol, indicates a hazardous situation which, if not avoided, could result in minor or moderate injury.

NOTICE

NOTICE is used to address practices not related to personal injury.

Note A note presents information that is important, but not hazard-related.

IMPORTANT An important notice contains information essential to software or hardware operation.

Typographical conventions

VMware Smart Assurance uses the following type style conventions in this document:

Bold	Use for names of interface elements
Italic	Use for full titles of publications referenced in text
`Monospace`	Use for: ■ System output, such as an error message or script ■ System code ■ Pathnames, filenames, prompts, and syntax ■ Commands and options

`Monospace italic`	Use for variables.
`Monospace bold`	Use for user input.
[ ]	Square brackets enclose optional values
|	Vertical bar indicates alternate selections — the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

6

VMware Smart Assurance support, product, and licensing information can be obtained as follows:

Product information

For documentation, release notes, software updates, or information about VMware Smart Assurance products, go to VMware Online Support at:

docs.vmware.com

Technical support

Go to VMware Smart Assurance Online Support. You will see several options for contacting VMware Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your VMware sales representative for details about obtaining a valid support agreement or with questions about your account.

Your comments

7

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Provide your opinions of this document to:

docs.vmware.com

Introduction

8

This chapter includes the following topics:

- Overview
- Topology representation and root-cause analysis
- Edge connections and proxy devices
- Architecture
- Installation
- Service command
- Configuration
- Discovery
- Invocation
- Operating modes
- Full and incremental topology splits
- Core topology-split functions
- Trial learning mode
- Automatic topology synchronization
- Edge management
- TSM hook script

Overview

The VMware Smart Assurance Topology Split Manager is a specially configured instance of an IP Manager that is started as a “DISCOVERY_MANAGER.” As such, the Topology Split Manager discovers availability-type (but not performance-type) topology objects and does no monitoring or analysis of the discovered topology.

The Topology Split Manager provides scalable topology management for IP network domains that contain a large number of devices. When run in normal mode with all operating features enabled, the Topology Split Manager performs the following tasks:

- Performs an initial discovery of the devices
- Splits the discovered devices across multiple domains that are managed by individual IP Availability Managers
- Creates a seed file of devices for each domain
- Places the devices that are listed in a domain's seed file onto the Pending Devices list of the domain's IP Availability Manager
- Adds proxy devices to the domains so that events from the IP Availability Managers can be stitched together correctly by the VMware Smart Assurance Service Assurance Manager (Global Manager) to point to the common root cause

  Given a large topology, like the one in Large IP network, the goal of the Topology Split Manager is to split the devices in the topology into a user-specified number of domains such that the domains are fairly equal in size and the number of proxy devices is kept to a minimum.

Figure 8-1. Large IP network



Splitting a large IP network into four domains—Example is an example of a large topology that has been split by the Topology Split Manager into four domains, each of which is managed by an IP Availability Manager.

Figure 8-2. Splitting a large IP network into four domains—Example



An “IP Availability Manager” may be either an IP Availability Manager (AM) or an IP Availability Manager and an IP Performance Manager running as a single process (AM-PM).

Topology representation and root-cause analysis

The Topology Split Manager splits the devices in the topology into multiple domains such that:

- The topology representation by the multiple IP Availability Managers is equivalent to the topology representation by a single IP Availability Manager *if that manager were able to manage all devices in the topology*.
- The root-cause analysis by the multiple IP Availability Managers is equivalent to the root-cause analysis by a single IP Availability Manager *if that manager were able to manage all devices in the topology*.

Edge connections and proxy devices

When topology is split into multiple domains, connections that cross domain boundaries are broken. To monitor such connections to ensure proper correlation across all domains, the Topology Split Manager adds proxy devices, or replicated devices, to neighboring domains. Representation of a connection that crosses a domain boundary—Example provides an example.

Figure 8-3. Representation of a connection that crosses a domain boundary—Example



A. Representation of the connections in the original topology



B. Representation of the same connections after the topology split

The domain that contains the “real” device in its topology is said to “own” the device and has primary responsibility for managing the device. As such, the domain is responsible for managing all components of the device.

The domain that contains the “proxy” device in its topology has partial management responsibility for that device. As such, the domain is responsible for managing only those components of the proxy device that are involved in the “across the domain” connection from the proxy device to the device that is owned by the domain (Router C in Figure 3).

The Topology Split Manager determines the proxy devices for the domains, adds them to the domain seed files, and places them on the appropriate IP Availability Manager’s Pending Devices list for subsequent pending discovery. During the discovery of the proxy devices, and assuming that an operating feature called *TSM hook script* is enabled, the IP Availability Manager sets the Domain attribute of the devices to “Proxy.”

Architecture

Topology-split architecture shows the Topology Split Manager architecture for a deployment that has <n> domains named AM001, . . ., and AM00<n>.

Figure 8-4. Topology-split architecture



The Topology Split Manager splits the topology into <n> domains, which reduces the scope of topology that is managed by each IP Availability Manager while minimizing the overlap of the discovered devices within the domains and ensuring that no coverage gaps exist across the domains.

Through the *optional* VMware Smart Assurance Adapter for Alcatel-Lucent 5620 SAM EMS, the Topology Split Manager receives the topology information that is needed by the IP Availability Managers to model networks that are populated with Alcatel-Lucent devices. The VMware Smart Assurance Adapter for Alcatel-Lucent 5620 SAM EMS User Guide describes the Adapter for Alcatel-Lucent 5620 SAM EMS in detail and presents configuration procedures for the Adapter.

After the topology split, the IP Availability Managers discover, monitor, and analyze the devices and network connectivity in their respective domains, and send topology and event information to the Global Manager. The Global Manager integrates the topology and event information to acquire the complete topology of the managed IP network domain and to diagnose problems and impacts for the managed IP network domain.

The Topology Split Manager is the central point for topology management. All topology changes are made through the Topology Split Manager rather than through the IP Availability Managers that manage the domains.

Installation

The Topology Split Manager software is included in the VMware Smart Assurance IP Manager. When installing the IP Manager, you will not see any mention of the Topology Split Manager.

Supported platforms for the Topology Split Manager

The Topology Split Manager software is supported on Sun Solaris and Red Hat Linux. For a list of supported OS versions for the IP Manager, see the VMware Smart Assurance Installation Guide for SAM, IP, ESM, MPLS, and NPM Managers.

Guidelines for installing the Topology Split Manager

Follow these guidelines to install the Topology Split Manager:

- Install the IP Manager to a new installation directory.

  Include TSM in the directory pathname for ease in locating this particular installation. For example, on UNIX:

  /opt/InCharge/TSM/IP

  In this document, the term TSM_BASEDIR represents the base directory, <installation directory>/TSM/IP, where the software for the Topology Split Manager is installed.

- Choose a “Complete” installation.

- Do not install any products as services.

  Detailed instructions for installing the IP Manager are given in the VMware Smart Assurance Installation Guide for SAM, IP, ESM, MPLS, and NPM Managers.

Installed file tree for the Topology Split Manager

The installed file tree for the Topology Split Manager is shown in Topology Split Manager file tree and described in Topology Split Manager files.

Figure 8-5. Topology Split Manager file tree



Table 8-1. Topology Split Manager files

Filename	Description
Files in rules/discovery/topo-split directory	
Topo-split scripts on Topology Split Manager: ■ topo-split-init.asl ■ topo-split-import.asl ■ topo-split.asl ■ topo-split-output-seeds.asl ■ topo-split-output-seeds-AM.asl ■ topo-split-report.asl ■ topo-split-trial-hook.asl ■ topo-split-sync.asl ■ topo-split-sync-edge.asl ■ topo-split-sync-AM.asl ■ topo-split-hook.asl ■ topo-split-override-polling-system.asl ■ topo-split-hook-AM.asl Topo-split scripts on IP Availability Manager: ■ ic-post-topo-management.asl ■ ic-override-polling-post-processor.asl ■ ic-ts-manage-proxy-post-processor.asl ■ ic-manage-hsrp-post-processor.asl	Scripts that the Topology Split Manager runs when performing a topology split. Which scripts run depends on which of the following features are enabled for the topology split: ■ Operating mode: normal or trial learning mode ■ Automatic topology synchronization ■ Edge management ■ TSM hook script The rules/discovery/topo-split directory also contains scripts that a domain's IP Availability Manager runs when automatic topology synchronization, edge management, and the TSM hook script are enabled.
Files in conf/discovery/topo-split directory	
tsm.conf	Configuration file that a user edits to control the topology-split operation.
topo-split.conf	Configuration file that a user edits to specify the names of the domains into which the topology will be split.
Files in conf/discovery directory	
discovery.conf tpmgr-param.conf name-resolver.conf partitions.conf user-defined-connections.conf	Configuration files that a user edits to control Topology Split Manager discovery.
Files in conf/icf directory	
bootstrap-tsm.conf	Bootstrap file for the Topology Split Manager.
Files in conf/custom directory	
override-polling-system.txt	Text file in which a user specifies a list of devices that the user wants the IP Availability Managers to poll at a one-minute interval.
Files in local/logs directory	
Server log file Example: INCHARGE_TSM.log	A server log in which startup, progression, and error messages are written for the Topology Split Manager.

Table 8-1. Topology Split Manager files (Continued)

Filename	Description
seedfile-<AM1 name>-<date>, ..., seedfile-<AMn>-<date> Example: seedfile-AM001-1138291391	Seed files that are created for a full or incremental topology split when the Topology Split Manager is operating in normal mode. Each seed file contains the split topology (devices) for an IP Availability Manager.
<TSM name>.trials.csv Example: INCHARGE-TSM.trials.csv	A comma-separated values (csv) report file that is created for a full topology split when the Topology Split Manager is operating in trial learning mode. For each successful full topology split during a trial learning mode invocation, the Topology Split Manager adds a row of data to the report file. Existing data is removed from the <TSM name>.trials.csv file at the beginning of each new invocation of trial learning mode.
Files in local/logs/<TSM name>.trial.<num> directory	
<TSM name>.rpsExample:INCHARGE-TSM.rps seedfile-<AM1 name>, ..., seedfile-<AMn>Example:seedfile-AM001 report.csv summary.csv	Repository snapshot file, seed files, and csv report files that are created for a full topology split when the Topology Split Manager is operating in trial learning mode. For each successful full topology split during a trial learning mode invocation, the Topology Split Manager creates a new set of files and a new <TSM name>.trial.<num> directory into which to store the files. Existing <TSM name>.trial.<num> directories are removed at the beginning of each new invocation of trial learning mode.
Files in local/repos/icf directory	
Repository file Example: INCHARGE-TSM.repos	A repository file in which the repository of the Topology Split Manager is stored.

License required to use the Topology Split Manager

The Topology Split Manager checks out the following IP Manager licenses at startup:

- AP_ALCT_SAM_USR
- AP_AUTO_DISCOVERY
- AP_DISCOVERY
- AP_HOST_RES
- AP_ICSRV
- AP_IPV6
- AP_L3S
- AP_NAS
- AP_VIRT_ROUTER
- AP_VRRP
- DMC_CORR

- DMC_DECS
- VN_ALL
- VN_CISCO
- VN_CISCO_HSRP

  Even though the Topology Split Manager does not use the DMC_CORR license or (in a default configuration) the DMC_DECS license, it does check out those licenses at startup.

  The Topology Split Manager does not require a volume license.

Service command

From the TSM_BASEDIR/smarts/bin directory in the IP Manager installation area, a user manually installs the Topology Split Manager as a service after the installation. Here is the Topology Split Manager service command.

UNIX:

```
/opt/InCharge/TSM/IP/smarts/bin/sm_service install
 --force
 --name=ic-tsm
 --startmode=runonce
 --description="VMware Smart Assurance Topology Split Manager"
 /opt/InCharge/TSM/IP/smarts/bin/sm_server
 --name=INCHARGE-TSM
 --config=icf
 --bootstrap=bootstrap-tsm.conf
 --port=0
 --ignore-restore-errors
 --noreconfigure
 --output
--daemon
```

To install the service, invoke the service command on one line.

Startup

To manually start the service and the Topology Split Manager, invoke the following command:

sm_service start ic-tsm

Shutdown

To manually stop the service and the Topology Split Manager, invoke the following command:

sm_service stop ic-tsm

Configuration

Before invoking a topology split, a user edits a file named topo-split.conf to specify the names of the domains into which the topology will be split, and edits a file named tsm.conf (or invokes dmctl commands) to control the following aspects of the topology split:

- Enable the topology split

  The associated configuration parameter is NumberOfDomains (Default: 1).

- Select an operating mode for the topology split: normal mode or trial learning mode

  The associated configuration parameters are:

  - IsTrialLearningEnabled (Default: FALSE)
  - NumberOfTrials (Default: 3)

- Configure operating features for the topology split

  The associated configuration parameters are:

  - IsSplitFromScratch (Default: TRUE)
  - IsTopoSyncEnabled (Default: FALSE)
  - EnableAutoEdgeManagement (Default: FALSE)
  - IsTSMHookEnabled (Default: FALSE)

- Fine tune the topology-split algorithm

  The associated configuration parameters are:

  - AdapterHighWaterMark (Default: 100000)
  - DeviceHighWaterMark (Default: 5000)

    .

    .

    .

  - NumberOfRetries (Default: 3)

    By default, topology split is disabled because the default value of the NumberOfDomains configuration parameter is 1. To enable topology split, a user must set NumberOfDomains to the number of domains that are specified in the topo-split.conf file. That number must be greater than 1.

    The tsm.conf file and the topology-split configuration parameters are described in *"Description of tsm.conf" on page 52*. The topo-split.conf file and the "DOMAINNAME" configuration parameter are described in *"Description of topo-split.conf" on page 59*.

    Chapter 9 Configuring the Topology Split Manager identifies the configuration tasks in a Topology Split Manager deployment and provides the procedures for completing those tasks.

Discovery

During the initial full discovery, the Topology Split Manager populates its repository with the devices and network connections that it discovers in the managed IP network domain. During any subsequent full discovery, the Topology Split Manager rediscovers the devices and network connections in the managed IP network domain.

The discovery flow up to postprocessing is essentially the same as for any IP Manager except that the discovery is more streamlined. During the postprocessing phase of discovery, and only if the NumberOfDomains configuration parameter is set, the Topology Split Manager starts a topology-split driver that invokes topology split.

Chapter 9 Configuring the Topology Split Manager identifies the discovery configuration tasks in a Topology Split Manager deployment and provides the procedures for completing those tasks. Chapter 10 Initiating Discovery and Topology Split presents procedures for initiating discovery in a Topology Split Manager deployment.

Invocation

After all components in a Topology Split Manager deployment have been installed and configured, a user starts all components in the deployment and initiates a full discovery. When per-device (phase 2) discovery completes, and assuming that the NumberOfDomains configuration parameter is set, the Topology Split Manager starts a topology split driver that invokes topology split.

In addition, if the Topology Split Manager has already performed a full discovery of the devices in the managed IP network domain (and, again, assuming that the NumberOfDomains configuration parameter is set), the user can enter the following command from the TSM_BASEDIR/smarts/bin directory to invoke topology split:

```
dmctl -s <TSM name> invoke GA_Driver::TopoSplit-Driver start
```

where <TSM name> is the name of the Topology Split Manager.

Operating modes

The Topology Split Manager can operate in one of two modes: normal mode or trial learning mode. The value of the IsTrialLearningEnabled configuration parameter determines the operating mode:

- When IsTrialLearningEnabled is FALSE (default), the Topology Split Manager runs in normal mode and creates a single topology split when topology split is invoked.
- When IsTrialLearningEnabled is TRUE, the Topology Split Manager runs in trial learning mode and creates multiple topology splits when topology split is invoked.

  Trial learning mode is described in Trial learning mode.

Full and incremental topology splits

When operating in normal mode and topology split is invoked, the Topology Split Manager performs either a *full* topology split or an *incremental* topology split. The value of the IsSplitFromScratch configuration parameter determines the type of topology split:

- When IsSplitFromScratch is TRUE (default), the Topology Split Manager performs a full topology split.

  For a full topology split, the Topology Split Manager splits the topology into as many domains as specified in the NumberOfDomains configuration parameter and then clears IsSplitFromScratch (sets IsSplitFromScratch to FALSE).

- When IsSplitFromScratch is FALSE, the Topology Split Manager performs an incremental topology split.

  For an incremental topology split, the Topology Split Manager applies incremental topology updates to an existing topology split, which might result in the moving of some devices from one domain to another in an attempt to improve the split.

  A full Topology split should be run only for an initial full discovery or after a major network redesign. Incremental topology split should be used for normal topology updates. Chapter 10 Initiating Discovery and Topology Split describes full and incremental topology splits in more detail.

Core topology-split functions

When topology split is invoked, the topology-split driver runs scripts that read the tsm.conf and topo-split.conf files. If the NumberOfDomains configuration parameter is set to the number of domains in the topo-split.conf file, additional scripts run, such as the ones that perform the following tasks:

- Run the topology-split algorithm, which analyzes and splits the discovered topology into the user-specified number of domains. The topology-split algorithm may be run more than once, depending on the value of the IsTrialLearningEnabled configuration parameter, and on the value of the NumberOfRetries configuration parameter if the topology split fails.

  Any retry of a failed topology split is performed as an incremental topology split.

- In normal mode (IsTrialLearningEnabled = FALSE), create a seed file for each domain and write the seed files to the Topology Split Manager's local/logs directory.
- In trial learning mode (IsTrialLearningEnabled = TRUE and IsSplitFromScratch = TRUE), create a seed file for each domain *for each successful run of the topology-split algorithm*, and write the seed files to a Topology Split Manager's <TSM name>.trial.<num> directory.
- For a discovery that includes Hot Standby Router Protocol (HSRP) groups and/or Virtual Router Redundancy Protocol (VRRP) groups, make corrections for any HSRP or VRRP group that is split between domains as a result of the topology split, to ensure that at least one of the domains includes the complete topology for an HSRP or VRRP group. All of the devices for a "complete" group are added to the domain that contains the lowest number of *managed* network adapters.

Topology seed files created in normal mode

In normal mode, the Topology Split Manager creates seed files in its local/logs directory for both full topology splits and incremental topology splits.

Each seed file for a full topology split contains the discovered devices for the associated IP Availability Manager. Each new seed file for an incremental topology split contains the previously discovered devices and the newly discovered devices for the associated IP Availability Manager. Also, devices (if any) that have been removed from the managed IP network domain will be removed from the new seed files for an incremental topology split.

The devices in a seed file are identified as either primary seeds or proxies. Primary seed devices, which are devices that are owned by the associated IP Availability Manager, appear in the “Primary seeds” section of the seed file. Proxy devices appear in the “Proxies” section of the seed file.

The name format of a seed file is:

```
seedfile-<AM name>-<date>
```

where:

- <AM name> is the domain name for which the seed file was generated. The name of the domain and the name of the IP Availability Manager that manages the domain are the same name.
- <date> is a timestamp that identifies when the seed file was generated, stated as the number of seconds since January 1, 1970, 00:00:00 GMT.

  For example, the following seed file is for the domain AM001:

  ```
  seedfile-AM001-1138291391
  ```

  When automatic topology synchronization is not enabled (default), users can use the seed files to initiate IP Availability Manager discovery manually.

Topology seed files created in trial learning mode

In trial learning mode, the Topology Split Manager creates seed files in its <TSM name>.trial.<num> directories for full topology splits only.

The name format of a seed file in a <TSM name>.trial.<num> directory is:

```
seedfile-<AM name>
```

where <AM name> is the domain name for which the seed file was generated.

For example, the following seed file is for the domain AM001:

```
seedfile-AM001
```

Upon choosing an initial topology split, users can use the seed files to initiate IP Availability Manager discovery manually.

Trial learning mode

Because the Topology Split Manager's topology-split algorithm is pseudo-random, different runs of the algorithm might produce different results. A new feature, called *trial learning mode*, enables users to run multiple initial (full) topology splits in an automated fashion, so that they can compare a number of initial splits and then choose the one that best matches their criteria.

Trial learning mode enablement

Operating in trial learning mode requires the following configuration parameter settings for the Topology Split Manager:

- NumberOfDomains = <*number of domains into which to split the topology*>
- IsSplitFromScratch = TRUE
- IsTrialLearningEnabled = TRUE
- NumberOfTrials = <*number of initial (full) topology splits to run*>

  When operating in trial learning mode, the Topology Split Manager ignores the following configuration parameters:

- IsTopoSyncEnabled
- EnableAutoEdgeManagement
- IsTSMHookEnabled

  That is, regardless of the settings of these configuration parameters, the Topology Split Manager will not run automatic topology synchronization, edge management, or the TSM hook script.

Trial learning mode operation

When operating in trial learning mode, the Topology Split Manager performs the following tasks:

- For each trial topology split, splits the topology into as many domains as specified in the NumberOfDomains configuration parameter.
- Upon creating the last trial topology split, clears the IsSplitFromScratch configuration parameter (sets IsSplitFromScratch to FALSE).

  For each successful trial topology split, the Topology Split Manager creates a directory named <TSM name>.trial.<num> in the TSM_BASEDIR/smarts/local/logs directory, as shown in Directories and files created by the Topology Split Manager in trial learning mode.

Figure 8-6. Directories and files created by the Topology Split Manager in trial learning mode



<TSM name> is the name of the Topology Split Manager, and <num> is a number that is padded with leading zeros to three digits and identifies a successful trial topology split. An example of <TSM name>.trial.<num> is INCHARGE-TSM.trial.001, where 001 identifies the first successful trial topology split.

The Topology Split Manager populates each <TSM name>.trial.<num> directory with the following files:

- A seed file for each domain's IP Availability Manager, named seedfile-<AM1 name>, ..., seedfile-<AMn name>; for example: seedfile-AM001

  Unlike a topology split in normal mode, in which the seed files are stored in the local/logs directory and have a date stamp in their names (seedfile-<AM name>-<date>), the seed files for a topology split in trial learning mode are stored in a <TSM name>.trial.<num> directory and do not have a date stamp in their names.

- A repository (rps) file named <TSM name>.rps; for example: INCHARGE-TSM.rps

  The repository file is a snapshot of the repository state for the topology split.

- Two comma-separated values (csv) report files named report.csv and summary.csv

  The Topology Split Manager also creates a csv report file named <TSM name>.trials.csv in the TSM_BASEDIR/smarts/local/logs directory and populates it with a row of data for each topology split.

Note Existing <TSM name>.trial.<num> directories are removed at the beginning of each new "trial learning mode" invocation.

Trial learning mode report files

For each successful trial topology split, the Topology Split Manager creates a <TSM name>.trial.<num> directory and populates the directory with the report.csv and summary.csv files. It also adds a line of summary data for the topology split to the <TSM name>.trials.csv file, which is created after the first successful topology split.

The report files contain criteria such as (1) the number of devices that are owned by a domain and (2) the number of devices that are proxied by a domain. For a device that is owned by a domain, the domain has the primary responsibility for managing the device. For a device that is proxied by a domain, the domain has a partial responsibility for managing the device. Partial management means that the domain is responsible for managing only those components of the proxy device that are involved in the "across the domain" connections from the proxy device to devices that are owned by the domain.

The report.csv file contains one row for each device, for each domain in which the device appears. The file has the following columns of information:

- Name of the domain
- Class of the device (Switch, Router, and so on)
- Name of the device
- Whether the device is owned by the domain, or proxied

- Number of ports
- Number of interfaces
- SNMP address
- SNMP System OID
- Description

summary.csv

The summary.csv file contains one row for each domain, and summarizes the distribution of topology between the domains. The file has the following columns of information:

- Name of the domain
- Number of devices that the domain owns

  Devices for which the domain has the primary responsibility for management.
- Number of ports on owned devices
- Number of interfaces on owned devices
- Topology Split Manager weight metric for the owned devices
- Number of devices that the domain proxies

  Devices for which the domain has a partial responsibility for management.
- Number of ports on proxied devices
- Number of interfaces on proxied devices
- Topology Split Manager weight metric for the proxied devices
- Total of the owned and proxied weight metrics
- Percentage relative to an equal share of the weights that this domain contains.

  That is, for a perfect split, each domain would carry exactly 100% of their fair share of the weight. Realistically, some domains will carry more than their fair share (>100%) and other will carry less (<100%). The closer these numbers get to 100% across all domains, the more evenly distributed the load becomes across the domains.

TSM name.trials.csv

This file contains one row of summary data for each successful topology split. (Terms such as “imbalance ratio” and “weight” are defined in #unique_47/unique_47_Connect_42__topo_split_configure_58059.) The file has the following columns of information:

- Number of the trial

- Flags related to this trial

  This column is a string column that contains zero or more of the following strings, separated by a slash (/) if there are more than one:

  - BEST_BALANCE if this trial is the most evenly balanced trial (where “most evenly balanced” equates to the lowest imbalance ratio)
  - LOWEST_PROXIES if this trial has the lowest proxy percentage (where “proxy percentage” is the total number of proxy devices in all domains *divided by* the total number of proxied and owned devices in the topology split)

    Note that more than one trial can have the same flag if the trials tie for the best value on that criterion.
- Minimum weight of any of the domains in the trial
- Maximum weight of any of the domains in the trial
- Average weight of the domains
- Imbalance ratio, specified as the ratio of the maximum to average weight

  This criterion might be used in evaluating the best split; lower is better.
- Percentage of the total number of devices across all domains that are proxies

  Another useful criterion in judging the best split; lower is better.
- Name of the domain with the highest number of ports and interfaces

  Used with the following two fields to evaluate whether the topology split exceeds scalability estimates for the domain’s IP Availability Manager.
- Number of ports in that domain
- Number of interfaces in that domain

Note Existing data is removed from the `TSM name.trials.csv` file at the beginning of each new invocation of trial learning mode.

Creating additional report files

To handle cases where the user has additional criteria to report, a script named topo-split-trial-hook.asl is invoked after each topology split. A placeholder version of this script is shipped with IP Manager 9.1, and is located in the TSM_BASEDIR/smarts/rules/discovery/topo-split directory.

The topo-split-trial-hook.asl has a TRIAL_HOOK section that can be filled in to perform additional logic after each topology split, between the generation of the standard reports and the time that the repository snapshot is taken. The script also has an END_HOOK section that is called once after all trials are run.

Particular techniques that might be useful include the use of ACT_Script or ACT_PerlScript or similar to run an external script, or ACT_File to generate data files directly from VMware's Adapter Scripting Language (ASL). For the latter case, the subdir property of TopoSplit_Manager could be used to initialize the subdir property of an ACT_File instance, to place the file in the appropriate trial directory.

Choosing an initial topology split

Users will compare the resulting seed files, csv reports, and trial-hook-created reports (if any) to choose the initial topology split that they prefer. Here are some guidelines for evaluating and choosing the best topology split:

- A split with a lower imbalance ratio is preferable to a split with a higher imbalance ratio.
- A split with a lower total proxy percentage is preferable to one with a higher total proxy percentage.

  The true load that a topology puts on a domain depends on other statistics as well.

  A user might want to experiment with different values of the topology-split algorithm fine-tuning parameters that are described in *"Description of tsm.conf" on page 52*, but should consider doing so only after running the topology split with the default values.

  After choosing a trial initial topology split, a user should perform the procedure in Using an initial topology split from trial learning mode to continue with the desired trial result as a starting point.

Automatic topology synchronization

At the end of a topology split and assuming that the IsTopoSyncEnabled configuration parameter is TRUE, the Topology Split Manager automatically distributes the split topology to the IP Availability Managers and synchronizes the discovered topologies of the IP Availability Managers with the assigned topologies in the Topology Split Manager.

Automatic topology synchronization enablement

Automatic topology synchronization requires the following configuration parameter settings for the Topology Split Manager:

- NumberOfDomains = <*number of domains into which to split the topology*>
- IsTopoSyncEnabled = TRUE
- IsTrialLearningEnabled = FALSE

  In normal mode, automatic topology synchronization can be enabled for a full topology split (IsSplitFromScratch = TRUE) or an incremental topology split (IsSplitFromScratch = FALSE). In trial learning mode, automatic topology synchronization does not run, regardless of the setting of the IsTopoSyncEnabled configuration parameter.

Automatic topology synchronization operation

When a topology split is invoked, the topology-split driver runs scripts that perform core topology-split tasks, as described in Core topology-split functions. If IsTopoSyncEnabled is TRUE (and IsTrialLearningEnabled is FALSE), additional scripts run, such as the ones that perform the following tasks:

- Create seed files in the IP Availability Managers *Create seed files in the IP Availability Managers*
- Synchronize topology with the IP Availability Managers *Synchronize topology with the IP Availability Managers*

Create seed files in the IP Availability Manager s

The Topology Split Manager forwards the list of devices in each seed file in the local/logs directory to the appropriate IP Availability Manager. There, under the control of the Topology Split Manager, the IP Availability Manager creates a seed file in its local/logs directory and populates the file with the list of devices from the Topology Split Manager.

The Topology Split Manager directs the IP Availability Manager to add the string “update” to the name of the seed file that it creates. For example, if the name of the source seed file in the Topology Split Manager is *seedfile-AM001-1138291391*, the IP Availability Manager will name the seed file *seedfile-* **update-** *AM001-1138291391*.

The Topology Split Manager obtains a list of discovered devices from the IP Availability Manager and then calculates the differences *between* the actual devices that were previously discovered by the IP Availability Manager *and* the devices that are assigned to the IP Availability Manager by the topology-split algorithm. The Topology Split Manager uses this information to organize the devices in the IP Availability Manager-created seed file, hereafter called the *domain’s seed file*, as “New Seeds,” “Unmanaged Seeds,” or “Outdated Seeds.”

“New Seeds” contains an entry for each *managed* device (IsManaged attribute = TRUE) that is assigned by the topology-split algorithm (as either an owned or proxy device) to the IP Availability Manager *and* is not already in the IP Availability Manager’s topology.

“Unmanaged Seeds” contains an entry for each *unmanaged* device (IsManaged attribute = FALSE) that is assigned by the topology-split algorithm (as either an owned or proxy device) to the IP Availability Manager. The entries in “Unmanaged Seeds” are commented-out.

“Outdated Seeds” contains an entry for each device (owned or proxy) that is already in the in the IP Availability Manager’s topology *and* is either not assigned by the topology-split algorithm or is assigned and is *unmanaged*. The entries in “Outdated Seeds” are commented-out.

The IP Availability Manager will discover only the devices in “New Seeds.”

If topology synchronization is successful, a user does not have to do anything with a domain’s seed file.

Synchronize topology with the IP Availability Manager s

The Topology Split Manager performs the following tasks to synchronize the topology of an IP Availability Manager with the topology that is assigned to the IP Availability Manager by the topology-split algorithm:

- Creates and populates a table with the devices that are listed in the “Add Seeds” section of the domain’s seed file, and then adds those devices to the Pending Devices list of the IP Availability Manager.
- Creates and populates a table with the devices that are listed in the “Outdated Seeds” section of the domain’s seed file, and then removes those devices and their instrumentation objects from the IP Availability Manager’s repository.

  In this release, the TopoSplit_Manager can no longer be configured to initiate IP Availability Manager pending discovery at the end of topology synchronization. The pending discovery will occur at the beginning of the next scheduled automatic pending discovery, or will occur when a user manually invokes pending discovery from the command line or the Global Console. See Chapter 10 Initiating Discovery and Topology Split for more information about initiating IP Availability Manager pending discovery.

Edge management

At the end of a topology split and assuming that the IsTopoSyncEnabled and EnableAutoEdgeManagement configuration parameters are both TRUE, the Topology Split Manager performs the following tasks:

- Identifies which network adapters (ports or interfaces) on a proxy device do *not* have a connection that is fully monitored by the domain.
- Directs the domain’s IP Availability Manager to unmanage those network adapters during a subsequent pending discovery.

  Edge management reduces management overhead by unmanaging network adapters that attach to incomplete connections.

  As an aid to understanding the discussions that follow, Data model representations of network adapters and their connections”Examples is presented to show how network adapters and their connections are represented by the Topology Split Manager and the IP Availability Manager.

Figure 8-7. Data model representations of network adapters and their connections"Examples



A. Data model representation of router-to-router connection



B. Data model representation of router-to-switch connection



C. Data model representation of switch-to-switch connection

Not all relationships are shown in the figure. For example, neither the HostedBy relationship from an IP object to its host device nor the LayeredOver relationship from an IP object to its MAC is shown.

Edge management enablement

Edge management requires the following configuration parameter settings for the Topology Split Manager:

- NumberOfDomains = <*number of domains into which to split the topology*>
- IsTopoSyncEnabled = TRUE
- EnableAutoEdgeManagement = TRUE
- IsTrialLearningEnabled = FALSE

  Edge management runs as part of automatic topology synchronization.

  In normal mode, edge management can be enabled for a full topology split (IsSplitFromScratch = TRUE) or an incremental topology split (IsSplitFromScratch = FALSE). In trial learning mode, edge management does not run, regardless of the setting of the EnableAutoEdgeManagement configuration parameter.

Edge management operation

When a topology split is invoked, the topology-split driver runs scripts that perform core topology-split tasks, as described in Core topology-split functions. If IsTopoSyncEnabled and EnableAutoEdgeManagement are both TRUE (and IsTrialLearningEnabled is FALSE), additional scripts run, such as the ones that perform the following tasks:

- For each domain, build a table that contains entries for network adapters on proxy devices that are associated with connections that are *not* fully managed by the specified domain.
- Forward the respective tables to the respective IP Availability Managers.

  In each domain, during the postprocessing phase of pending discovery, the domain's IP Availability Manager performs the following tasks:

- Sets a flag on any protocol endpoint (IP or MAC, for example) that is LayeredOver any network adapter that has an entry in the table.

  The flag is a new attribute, named ForceLocalManagement, for protocol endpoints. ForceLocalManagement is FALSE by default.

- Unmanages each network adapter that has an entry in the table.

  Because of the set ForceLocalManagement flag, the unmanaging of the network adapters does *not* unmanage any protocol endpoints that are LayeredOver the network adapters. Without the flag, the unmanaging of a network adapter will unmanage any protocol endpoint that is LayeredOver the network adapter.

The IP Availability Manager keeps a list of protocol endpoints for which it has set the ForceLocalManagement flag. After all network adapters that have entries in the table have been unmanaged, the IP Availability Manager checks all instances of protocol endpoints and clears the flag on any protocol endpoint that is *not* in the list. This behavior ensures that ForceLocalManagement is FALSE on any protocol endpoint that might have had ForceLocalManagement set during an earlier topology split but no longer needs it set.

TSM hook script

At the end of a topology split and assuming that the IsTSMHookEnabled configuration parameter is TRUE, the Topology Split Manager runs the TSM hook script. The TSM hook script is a mechanism by which one-minute polling, proxy device management, and HSRP/VRRPGroup management information is transferred from the Topology Split Manager to the IP Availability Managers.

TSM hook script enablement

The TSM hook script requires the following configuration parameter settings for the Topology Split Manager:

- NumberOfDomains = <*number of domains into which to split the topology*>
- IsTSMHookEnabled = TRUE
- IsTrialLearningEnabled = FALSE

  In normal mode, the TSM hook script can be enabled for a full topology split (IsSplitFromScratch = TRUE) or an incremental topology split (IsSplitFromScratch = FALSE). In trial learning mode, the TSM hook script does not run, regardless of the setting of the IsTSMHookEnabled configuration parameter.

TSM hook script operation

When a topology split is invoked, the topology-split driver runs scripts that perform core topology-split tasks, as described in Core topology-split functions. If IsTSMHookEnabled is TRUE (and IsTrialLearningEnabled is FALSE), additional scripts run, such as the ones that run the following management code:

- One-minute polling management code *“One-minute polling management code”*
- Proxy device management code *“Proxy device management code”*
- HSRP/VRRPGroup management code *HSRP/VRRPGroup management code*

One-minute polling management code

This code identifies which devices are eligible for monitoring at a one-minute interval by reading the list of user-specified devices in the TSM_BASEDIR/smarts/conf/discovery/ custom/override-polling-system.txt file. (*“Description of override-polling-system.txt” on page 60* identifies the format of device entries in the override-polling-system.txt file.) The code then determines which devices in the list of devices belong to which domain, adds the devices to a table, and forwards the table to the IP Availability Managers.

Each IP Availability Manager reads the list of devices in the table.

For each device that is marked as belonging to a specific domain, the associated IP Availability Manager finds the component objects for the device and assigns a one-minute polling interval to the instrumented component objects of each device during the postprocessing phase of pending discovery. Exceptions are the following component object types:

- Processor
- Memory
- PowerSupply
- VoltageSensor
- TemperatureSensor
- Fan
- Network Adapter that is less than 34 megabytes (MB)

Note These objects are performance-type objects and are relevant to an IP Availability Manager and an IP Performance Manager running as a single process (AM-PM).

Proxy device management code

This code identifies which devices are proxy devices, adds the devices to a table, and forwards the table to the IP Availability Managers.

Each IP Availability Manager searches its topology for devices that are listed in the table. If a device is identified as a proxy in the table, the IP Availability Manager sets the device's Domain attribute to "Proxy" during the postprocessing phase of pending discovery.

Because the real device that is associated with a proxy device is already monitored for performance in another domain, the proxy device should not be monitored for performance in this domain. The proxy group is already created automatically in the Polling and Thresholds window of the IP domain manager. The identified proxy devices are now the member of the Proxy group under the Polling and Thresholds window.

Note This last statement is relevant to an IP Availability Manager and an IP Performance Manager running as a single process (AM-PM).

Regardless of whether a proxy device is monitored for performance, it is always monitored for availability. Specifically, the domain to which the proxy device has been assigned monitors each proxy device's network adapter that is involved in a connection to a device that is owned by that domain.

HSRP/VRRPGroup management code

This code finds every HSRPGroup and VRRPGroup object in the Topology Split Manager's topology and inserts the name and group size of each found group object into a table. It then forwards the table to the IP Availability Managers.

Each IP Availability Manager finds every HSRPGroup and VRRPGroup in its topology and compares its size to the size of the same group in the table. If the size of a HSRPGroup or VRRPGroup in the IP Availability Manager is smaller than the size of the same group in the Topology Split Manager, the IP Availability Manager unmanages the group during the postprocessing phase of pending discovery. Reason: A smaller group indicates an incomplete group, and an incomplete group might cause erroneous events.

Configuring the Topology Split Manager

9

This chapter includes the following topics:

- Overview
- Configuration directory
- User configuration parameters
- User configuration synchronization across installations
- Methods for modifying user configuration parameters
- Command sm_tpmgr and example
- Description of tsm.conf
- Description of topo-split.conf
- Description of override-polling-system.txt
- Tuning Topology Split Manager discovery
- Specifying domain names for topology split
- Specifying a list of devices for one-minute polling
- Enabling topology split
- Enabling trial learning mode
- Enabling automatic topology synchronization
- Enabling edge management
- Enabling the TSM hook script
- Configuring security

Overview

After installing the Topology Split Manager, the IP Availability Managers, the Global Manager, and the optional Adapter for Alcatel-Lucent 5620 SAM EMS in a Topology Split Manager deployment, you configure them so that they can form the proper connections to one another. You also configure discovery for the Topology Split Manager and the IP Availability Managers, and configure topology split for the Topology Split Manager.

Some configuration tasks are performed before the applications are started, such as the editing of configuration files, and some are performed through the Global Console when the applications are up and running.

Overview of Topology Split Manager configuration tasks

Configuring the Topology Split Manager consists of the following tasks:

- Copy the discovery configurations from one of the IP Availability Manager installations to the Topology Split Manager installation, and then modify certain discovery configurations in order to improve and streamline Topology Split Manager discovery. The discovery configurations are listed in Overview of Topology Split Manager configuration tasks.
- Optional: Enable autodiscovery.
- Enable and customize topology split.

  Procedures for configuring all aspects of IP Manager discovery are given in the VMware Smart Assurance IP Manager User Guide and the VMware Smart Assurance IP Manager Reference Guide.

  Procedures for improving Topology Split Manager discovery and configuring topology split are given in this chapter.

Overview of IP Availability Manager configuration tasks

In addition to performing the configuration and administration tasks that are common to all IP Availability Managers, you perform the following additional tasks to set up the IP Availability Managers in a Topology Split Manager deployment:

- Set the discovery configurations for one of the IP Availability Manager installations and then copy those configurations to the remaining IP Availability Manager installations.
- Disable autodiscovery.

  Procedures for configuring all aspects of IP Availability Manager are given in the VMware Smart Assurance IP Manager User Guide and the VMware Smart Assurance IP Manager Reference Guide.

Overview of Global Manager configuration tasks

In addition to performing the configuration and administration tasks that are common to all Global Managers, you perform the following additional task to set up the Global Manager in a Topology Split Manager deployment:

Use the Global Console and the Global Manager to specify the underlying IP Availability Managers.

Doing so associates the IP Availability Managers with the proper data exchange adapter (DXA) files, which identify the topology and events that the Global Manager should import from the IP Availability Managers.

Note No DXA or any other connection exists between the Topology Split Manager and the Global Manager.

Procedures for configuring all aspects of the Global Manager are given in the VMware Smart Assurance Service Assurance Manager Configuration Guide.

Overview of Adapter for Alcatel-Lucent 5620 SAM EMS configuration tasks

In addition to performing the configuration and administration tasks that are common to all Adapter for Alcatel-Lucent 5620 SAM EMS (Adapter) instances, you perform the following additional task to set up the Adapter in a Topology Split Manager deployment:

Set the AMServerName parameter value in the Adapter's emsConfig.import file to the name of the Topology Split Manager.

Upon reading this name, the Adapter will detect a Topology Split Manager destination and begin forwarding the topology data expected by a Topology Split Manager.

As an aside, the Topology Split Manager will set the ServiceName attribute of any topology object that is received from the Adapter to the name of the Adapter, to identify the Adapter as the object's source. When an Adapter-originating device object is assigned by the Topology Split Manager to an IP Availability Manager, the IP Availability Manager will use the ServiceName value to subscribe to the Adapter for status updates.

Procedures for configuring all aspects of the Adapter are given in the VMware Smart Assurance Adapter for Alcatel-Lucent 5620 SAM EMS User Guide.

Configuration directory

The Topology Split Manager configuration directory is TSM_BASEDIR/smarts/conf, shown in Configuration directory. TSM_BASEDIR, which is not an environment variable, is used in this document to represent the top-level directory structure of a Topology Split Manager installation. TSM_BASEDIR represents <installation root directory>/TSM/IP.

In the TSM_BASEDIR/smarts/conf directory:

- The "discovery" directory holds the discovery configurations for the Topology Split Manager.
- The "discovery/topo-split" directory holds the topology-split configurations for the Topology Split Manager.
- The "icf" directory holds the bootstrap file for the Topology Split Manager.

  The bootstrap file for the Topology Split Manager determines the operating role of the Topology Split Manager and sets up its environment.

The --config and --bootstrap options in a Topology Split Manager startup command point to the bootstrap-tsm.conf file in the icf directory. Pointing to this bootstrap file causes the startup software to load the bootstrap file, which contains a line of code that starts the Topology Split Manager as a "DISCOVERY_MANAGER."

Note The --config option in an IP Manager startup command points to the bootstrap files in the icf directory. The default bootstrap file is bootstrap.conf, which starts an IP Manager as an "AVAILABILITY_AND_PERFORMANCE_MANAGER" (AM-PM). The addition of the --bootstrap option in an IP Manager startup command points to a particular bootstrap file in the icf directory, such as bootstrap-am.conf, which starts an IP Manager as an "AVAILABILITY_MANAGER" (AM), or bootstrap-am-pm.conf, which also starts an IP Manager as an "AVAILABILITY_AND_PERFORMANCE_MANAGER" (AM-PM).

User configuration parameters

A major part of configuring the Topology Split Manager involves setting user configuration parameters in the files that are identified in User configuration files for the Topology Split Manager.

Table 9-1. User configuration files for the Topology Split Manager

Directory under TSM_BASEDIR/smarts	Filename	Description
conf/discovery	discovery.conf	File in which you fine tune configurations to improve the performance of Topology Split Manager discovery. See the VMware Smart Assurance IP Manager Reference Guide for a description of this file.
	tpmgr-param.conf	File in which you customize Topology Split Manager discovery. See the VMware Smart Assurance IP Manager Reference Guide for a description of this file.
	name-resolver.conf	File in which you control how the Topology Split Manager names a discovered system. See the VMware Smart Assurance IP Manager Reference Guide for a description of this file.
	partition.conf	File in which you specify display names for partitions. See the VMware Smart Assurance IP Manager Reference Guide for a description of this file.
	user-defined-connections.conf	File in which you specify connections between discovered network objects. See the VMware Smart Assurance IP Manager Reference Guide for a description of this file.
conf/discovery/topo-split	tsm.conf	File in which you enable and customize topology split. See *"Description of tsm.conf" on page 52* for a description of this file.
	topo-split.conf	File in which you specify domain names for topology split. See *"Description of topo-split.conf" on page 59* for a description of this file.

Table 9-1. User configuration files for the Topology Split Manager (Continued)

Directory under TSM_BASEDIR/smarts	Filename	Description
conf/custom	override-polling-system.txt	File in which you specify devices that the IP Availability Managers should poll at a one-minute interval. See *"Description of override-polling-system.txt" on page 60* for a description of this file.
local/conf	runcmd_env.sh	File in which you set environment variables to control behavior globally. For each application that is started from the Topology Split Manager installation area, the startup software sets the variables that are defined in the runcmd_env.sh file. The procedures that follow assume that the SM_BROKER or SM_BROKER_DEFAULT variable is defined in the runcmd_env.sh file and points to the Broker in the Topology Split Manager deployment.

As with most VMware Smart Assurance configuration files, the user configuration files for the Topology Split Manager are used to set attributes for certain objects within a Manager's environment when the Manager starts up. The parameters and their values in a configuration file become the attributes and their values for the created objects.

For the discovery.conf file, as an example, the associated object is an instance of the ICF_TopologyManager class and is named ICF-TopologyManager. This object has attributes for configuring various aspects of IP Manager discovery.

For the tsm.conf file, as another example, the associated object is an instance of the TopoSplit_Manager class and is named TopoSplit-Manager. This object has attributes for configuring various aspects of topology split.

User configuration synchronization across installations

Because the Topology Split Manager and the IP Availability Managers in a Topology Split Manager deployment may be started from different installations and run on different host machines, you will need to synchronize the discovery configuration parameters across the installations that are involved in the deployment. The synchronization involves the parameters in the following files in the TSM_BASEDIR or BASEDIR/smarts/conf/discovery directory:

- discovery.conf
- tpmgr-param.conf
- name-resolver.conf
- partition.conf

- user-defined-connections.conf

  As an example, the parameters in the name-resolver.conf file should be configured the same (NameFormat = “TM_USEAUTONAME” or NameFormat = “TM_USESEEDNAME”) across all installations. If local exceptions are defined, the same local exceptions should be defined for all installations.

  Note Due to differences in networking configuration on different host machines, VMware recommends that you set NameFormat to TM_USESEEDNAME on all installations. Doing so will ensure that the IP Availability Managers use the device names that are provided by the Topology Split Manager.

  In the Topology Split Manager installation, some of the discovery configurations in discovery.conf and tpmgr-param.conf should be modified in order to improve and streamline Topology Split Manager discovery, as described in *“Tuning Topology Split Manager discovery” on page 60*.

Methods for modifying user configuration parameters

There are three methods to modify user configuration parameters:

1. Edit the user configuration files and modify configuration parameters.
2. Issue the dmctl command to modify configuration parameters when the Topology Split Manager is running.
3. Use the Global Console to modify configuration parameters when the Topology Split Manager is running.

   Parameters that are modified by using methods 2 and 3 will be reconfigured to the settings in the user configuration files when the Topology Split Manager is shut down and restarted. Every time that the Topology Split Manager starts up, it sets its configurations in accordance to the settings in the user (and other) configuration files.

Editing configuration files to modify configuration parameters

VMware Smart Assurance provides the sm_edit utility to ensure that modified files are always saved to the appropriate local area, and that base (original) copies of the files remain unchanged. Preview of how the sm_edit utility works previews the operation of the sm_edit utility.

Figure 9-1. Preview of how the sm_edit utility works



All VMware Smart Assurance products use the same basic installation directory structure that is shown in Figure 8. For the IP Manager, from which the Topology Split Manager is launched, the product directory is named IP.

Note For a Topology Split Manager installation, TSM_BASEDIR represents <installation root directory>/TSM/IP.

To invoke the sm_edit utility, go to the TSM_BASEDIR/smarts/bin directory and specify the path and the name of the file, relative to the TSM_BASEDIR/smarts/local or TSM_BASEDIR/smarts directory, that you want to edit. For example,

```
sm_edit conf/discovery/topo-split/tsm.conf
```

opens in a text editor *either* a local copy of the tsm.conf file in TSM_BASEDIR/smarts/local/conf/discovery/topo-split *or* a base (original) copy of the tsm.conf file in TSM_BASEDIR/smarts/conf/discovery/topo-split if no local copy exists.

After you modify and save the tsm.conf file, the sm_edit utility saves the modified version of the file to the TSM_BASEDIR/smarts/local/conf/discovery/topo-split directory.

Chapter 12 The sm_edit Utility provides additional information.

Issuing the dmctl command to modify configuration parameters

You can use the following dmctl command to *view* the value of any basic-type attribute (string, boolean, integer, float, and so on) of any VMware Smart Assurance object:

```
dmctl -s <
   VMware Smart Assurance

 application name> get <class>::<instance>::<attribute>
```

You can use the following dmctl command to *change* the value of any basic-type attribute (string, boolean, integer, float, and so on) of any VMware Smart Assurance object:

```
dmctl -s <
   VMware Smart Assurance

 application name> put <class>::<instance>::<attribute> <attribute_value>
```

For the Topology Split Manager, as an example, you could use the dmctl **put** command to change attribute values for the ICF_TopologyManager object or the TopoSplit_Manager object.

Upon issuing a dmctl command, you might be prompted for a username and password. Respond with your user account. For example, to specify the default administrative account, enter username *admin* and password *changeme*.

The dmctl utility is described in the HTML pages that are located in the TSM_BASEDIR or BASEDIR/smarts/doc/html/usage directory of any VMware Smart Assurance product installation area.

Example 1

As an example of using the dmctl command to *view* an object's attribute value, enter the following dmctl command from the TSM_BASEDIR/smarts/bin directory of the Topology Split Manager to view the numberProbeThreads attribute value of the *ICF_TopologyManager* object:

```
dmctl -s <TSM name> get ICF_TopologyManager::ICF-TopologyManager::

numberProbeThreads
```

The default value for numberProbeThreads is 10.

To view all attribute values of the ICF_TopologyManager object in one invocation, enter the following dmctl command:

```
dmctl -s <TSM name> get ICF_TopologyManager::ICF-TopologyManager | more
```

Example 2

As another example of using the dmctl command to *view* an object's attribute value, enter the following dmctl command from the TSM_BASEDIR/smarts/bin directory of the Topology Split Manager to view the IsTopoSyncEnabled attribute value of the *TopoSplit_Manager* object:

```
dmctl -s <TSM name> get TopoSplit_Manager::TopoSplit-Manager::

IsTopoSyncEnabled
```

The default value for IsTopoSyncEnabled is FALSE.

To view all attribute values of the TopoSplit_Manager object in one invocation, enter the following dmctl command:

```
dmctl -s <TSM name> get TopoSplit_Manager::TopoSplit-Manager | more
```

Example 3

As an example of using the dmctl command to *change* an object's attribute value, enter the following dmctl command from the BASEDIR/smarts/bin directory of the Topology Split Manager to change the numberProbeThreads attribute value of the ICF_TopologyManager object:

```
dmctl -s <TSM name> put ICF_TopologyManager::ICF-TopologyManager::

numberProbeThreads 12
```

Your change will take effect immediately.

To make your change permanent, use the sm_edit utility to modify the numberProbeThreads parameter value in the conf/discovery/discovery.conf file.

Example 4

As another example of using the dmctl command to *change* an object's attribute value, enter the following dmctl command from the BASEDIR/smarts/bin directory of the Topology Split Manager to change the IsTopoSyncEnabled attribute value of the TopoSplit_Manager object:

```
dmctl -s <TSM name> put TopoSplit_Manager::TopoSplit-Manager::

   IsTopoSyncEnabled TRUE
```

Your change will take effect immediately.

To make your change permanent, use the sm_edit utility to modify the IsTopoSyncEnabled parameter value in the conf/discovery/topo-split/tsm.conf file.

Using the Global Console to modify configuration parameters

You can use the Global Console to view or change the attributes of any VMware Smart Assurance object. For the Topology Split Manager, as an example, you could use the Global Console to change attribute values for the TopoSplit_Manager object.

To use the Global Console to change attribute values for the TopoSplit_Manager object:

1. Attach the Global Console to the Topology Split Manager.

   The VMware Smart Assurance Service Assurance Manager Operator Guide provides information about starting the Global Console, logging in, and attaching to an VMware Smart Assurance application.

2. Press the **Ctrl-Alt-M** keys simultaneously to make all Topology Split Manager classes available to the Global Console, including the meta classes.

3. In the topology tree panel of the Global Console, right-click an open space and select **Refresh tree** from the pop-up menu to refresh the content of the topology tree.

4. In the topology tree panel, click the **+** signs to expand the **TopoSplit_Manager** class, then click **TopoSplit-Manager** to display the **Attributes** tab in the right panel.

   The **Attributes** tab lists the attribute names, data types, and values for an instance of the **TopoSplit_Manager** class named **TopoSplit-Manager**.

5. In the **Attributes** tab, use the scroll bar to scroll to the attribute that you want to change (**IsTopoSyncEnabled**, for example), double-click its value in the **Value** column, change the value, and press **Enter**.

6. Click **Apply** to save your changes.

7 Press the **Ctrl-Alt-M** keys simultaneously to make only the default Topology Split Manager classes available.

8 In the topology tree panel, right-click an open space and select **Refresh tree** from the pop-up menu to refresh the content of the topology tree.

9 In the topology tree panel, click the **+** sign next to the Topology Split Manager's name to expand the topology tree. Only the default classes appear in the topology tree.

   Your change will take effect immediately.

   To make your change permanent, use the sm_edit utility to modify the IsTopoSyncEnabled parameter value in the conf/discovery/topo-split/tsm.conf file.

Limitations

You can use the dmctl **put** command or the Global Console to override parameter values in the discovery.conf and name-resolver.conf files because:

- The parameters in these files are attributes of the ICF_TopologyManager object.
- The parameters in these files are loaded into memory at server startup.
- Parameter changes through dmctl occur *after* server startup and thereby override the corresponding parameter values in memory.

  However, you cannot use the dmctl **put** command or the Global Console to override parameter values in the tpmgr-param.conf file *because the parameters in that file correspond to entries for a table that is defined inside the ICF_TopologyManager object*. Instead, use the following dmctl **invoke** command to override parameter values in the tpmgr-param.conf file:

```
dmctl -s <TSM name> invoke ICF_TopologyManager::ICF-TopologyManager insertParameter
<parameter name> <value>
```

  For example:

```
dmctl -s <TSM name> invoke ICF_TopologyManager::ICF-TopologyManager insertParameter
DisableMPLSDiscovery TRUE
```

  You can also use the dmctl **put** command or the Global Console to override parameter values in the tsm.conf file because the parameters in this file are attributes of the TopoSplit_Manager object. However, because the parameters in the tsm.conf file are loaded into memory at topology-split invocation, *not* at server startup, dmctl-configured parameter changes are overridden when the tsm.file is loaded. Thus, using dmctl to override parameter values in the tsm.conf file does not work. What does work is to use dmctl to change the values of TopoSplit_Manager object attributes that are either *not* in the tsm.conf file or are commented-out in the tsm.conf file.

  And lastly, because the content of the partition.conf, user-defined-connections.conf, topo-split.conf, override-polling-system.txt, and runcmd_env.sh files are not attributes of any object, you must use the sm_edit utility to modify these files.

Command sm_tpmgr and example

You can use the sm_tpmgr command to reload the parameters in the discovery.conf, tpmgr-param.conf, or name-resolver.conf file. A discovery cycle should not be in progress when you issue this command.

After you use the sm_edit utility to make a parameter change in the discovery.conf, tpmgr-param.conf, or name-resolver.conf file, you can use the sm_tpmgr command to make that change known to the Topology Split Manager without the need to restart the Topology Split Manager.

The syntax for the sm_tpmgr command is:

```
sm_tpmgr -s <TSM name> --load-conf=<configuration filename>
```

For example:

```
sm_tpmgr -s INCHARGE-TSM --load-conf=discovery.conf
```

Description of tsm.conf

The tsm.conf file contains configuration parameters that you edit to enable and customize topology split. The parameters are a subset of attributes that are defined for the TopoSplit_Manager object.

When topology split is invoked, the Topology Split Manager loads the tsm.conf file.

The tsm.conf file contains the following active parameters:

```
NumberOfDomains = 1
IsTSMHookEnabled = FALSE
AdapterHighWaterMark = 100000
BalanceThreshold = 2.0000
DeviceHighWaterMark = 5000
NumberOfRetries = 3
IsDebugEnabled = FALSE
ProxyThreshold = 0.2500
```

It also contains the following commented-out parameters:

```
# EnableAutoEdgeManagement = TRUE
# IsTopoSyncEnabled = TRUE
# IsTrialLearningEnabled = FALSE
# IsSplitFromScratch = FALSE
# NumberOfTrials = 3
# IsSingleRootCauseEnabled = FALSE
# IsPreAssignmentEnabled = TRUE
# IsNonRelayDeviceFirstEnabled = TRUE
# MaxNumberOfRefinementLoops = 3
```

The parameters and their values are case-sensitive.

You can configure the topology-split parameters in the tsm.conf file, or you can configure them by using dmctl commands *as long as the configured parameter settings are not countermanded in tsm.conf*. Because the Topology Split Manager loads into memory whatever parameters are specified in tsm.conf at the beginning of a topology split, those settings will override any dmctl-configured settings for those parameters.

Attributes for TopoSplit_Manager::TopoSplit-Manager lists some key attributes for TopoSplit_Manager.

Table 9-2. Attributes for TopoSplit_Manager::TopoSplit-Manager

Attribute	Value	Description
topology-split mandatory parameters		
NumberOfDomains	Integer Default: 1	Number of domains that are registered with the Topology Split Manager. Must be greater than 1 and equal to the number of domains that are specified in the TSM_BASEDIR/smarts/local/conf/discovery/ topo-split/topo-split.conf file.
topology-split mode-of-operation attributes		
IsTrialLearningEnabled	TRUE,FALSE Default: FALSE	Determines whether the Topology Split Manager runs in normal mode or trial learning mode: ■ A value of TRUE indicates that the Topology Split Manager runs in trial learning mode and creates multiple full topology splits. ■ A value of FALSE indicates that the Topology Split Manager runs in normal mode and creates a single topology split. In normal mode, the IsSplitFromScratch attribute may be TRUE (to create a single *full* topology split) or FALSE (to create a single *incremental* topology split). In trial learning mode, the IsSplitFromScratch attribute *must* be TRUE. Also, in trial learning mode, the Topology Split Manager ignores the following attributes: ■ IsTopoSyncEnabled ■ EnableAutoEdgeManagement ■ IsTSMHookEnabled That is, regardless of the settings of these attributes, the Topology Split Manager will not run automatic topology synchronization, edge management, or the TSM hook script.
NumberOfTrials	Integer Default: 3	Number of initial (full) topology splits to run when the Topology Split Manager is running in trial learning mode. Must be greater than 1 when IsTrialLearningEnabled is TRUE.
topology-split operating feature attributes		

Table 9-2. Attributes for TopoSplit_Manager::TopoSplit-Manager (Continued)

Attribute	Value	Description
IsSplitFromScratch	TRUE,FALSE Default: TRUE	Determines whether the Topology Split Manager performs a full topology split or an incremental topology split: ■ A value of TRUE indicates that the Topology Split Manager performs a full topology split: Splits the topology into as many domains as specified in the NumberOfDomains attribute and then clears the IsSplitFromScratch attribute (sets IsSplitFromScratch to FALSE). ■ A value of FALSE indicates that the Topology Split Manager performs an incremental topology split: Applies incremental topology updates to an existing topology split, which might result in the moving of some devices from one domain to another in an attempt to improve the split. **Note** After creating a full topology split in normal mode or creating multiple full topology splits in trial learning mode, the Topology Split Manager automatically clears the IsSplitFromScratch attribute. In general, do *not* set IsSplitFromScratch in the tsm.conf file. (Instead, if need be, use dmctl to set this attribute.) That way, the reading of the tsm.conf file at the beginning of topology split will not override the automatic clearing of IsSplitFromScratch after full topology split.
IsTopoSyncEnabled	TRUE,FALSE Default: FALSE	Determines at the end of a topology split whether the Topology Split Manager performs an automatic topology synchronization between the Topology Split Manager and the domains' IP Availability Managers: ■ A value of TRUE indicates that the Topology Split Manager performs an automatic topology synchronization. ■ A value of FALSE indicates that the Topology Split Manager does not perform an automatic topology synchronization.
EnableAutoEdgeManagement	TRUE,FALSE Default: FALSE	Determines at the end of a topology split whether the Topology Split Manager performs edge management: ■ A value of TRUE indicates that the Topology Split Manager performs edge management. ■ A value of FALSE indicates that the Topology Split Manager does not perform edge management. Because edge management runs as part of topology synchronization, IsTopoSyncEnabled and EnableAutoEdgeManagement must both be TRUE for edge management to run.

Table 9-2. Attributes for TopoSplit_Manager::TopoSplit-Manager (Continued)

Attribute	Value	Description
IsTSMHookEnabled	TRUE,FALSE Default: FALSE	Determines at the end of a topology split whether the Topology Split Manager runs the TSM hook script: ■ A value of TRUE indicates that the Topology Split Manager runs the TSM hook script. ■ A value of FALSE indicates that the Topology Split Manager does not run the TSM hook script.
topology-split algorithm fine-tuning attributes		
AdapterHighWaterMark	Integer Default: 100000	Maximum number of network adapters (ports, interfaces) that should be placed into a single domain.
DeviceHighWaterMark	Integer Default: 5000	Maximum number of devices that should be placed into a single domain.
ProxyThreshold	Floating point number Default: 0.2500	Maximum ratio of total proxy devices in all domains *to* total number of proxied and owned devices in the Topology Split Manager.
ManagedAdapterWeight	Integer Default: 1	Weight for managed network adapters. Adjust the weights of ManagedAdapterWeight, UnmanagedAdapterWeight, and PerfAdapterWeight, *whose weights are relative to one another*, if you are not satisfied with the distribution of one type of network adapter (managed, unmanaged, or performance-managed) across domains. Because the default weight of ManagedAdapterWeight is the only non-zero weight of the three “weight” attributes, increasing or decreasing the weight of ManagedAdapterWeight and redoing the topology split will not change the split all that much. What will change the split (and result in a more balanced workload for the domains) is to increase the weight of UnmanagedAdapterWeight or PerfAdapterWeight, as explained in the descriptions for UnmanagedAdapterWeight and PerfAdapterWeight that follow.

Table 9-2. Attributes for TopoSplit_Manager::TopoSplit-Manager (Continued)

Attribute	Value	Description
UnmanagedAdapterWeight	Integer Default: 0	Weight for unmanaged network adapters. This weight is relative to the weight for ManagedAdapterWeight. Increase the weight of UnmanagedAdapterWeight to balance the allocation of unmanaged network adapters across the domains. If, for example, the topology-split algorithm allocates a lot of devices with unmanaged network adapters to one domain, which results in a longer discovery time for that domain relative to the other domains, you could increase the value of UnmanagedAdapterWeight (but to a value less than the value of ManagedAdapterWeight) and redo the topology split. The new split will *either* assign devices with unmanaged network adapters more evenly across the domains *or* include less devices with managed or performance-managed network adapters in the domain that has a lot of devices with unmanaged network adapters. Either way, the discovery time across the domains will be more balanced. **Note** Both monitoring time and discovery time for a domain increase with additional *managed* network adapters. Although monitoring time does not increase with additional *unmanaged* network adapters, the discovery time for a domain does increase. Balancing domains for discovery time could reduce how well domains are balanced for monitoring time.
PerfAdapterWeight	Integer Default: 0	Weight for network adapters of devices that are to have their components polled every minute. This weight is added to the weight of ManagedAdapterWeight to compute the weight for performance-managed network adapters. Increase the weight of PerfAdapterWeight to balance the allocation of performance-managed network adapters across the domains. If, for example, the topology-split algorithm allocates a lot of devices with performance-managed network adapters to one domain, which results in a higher workload for that domain relative to the other domains, you could increase the value of PerfAdapterWeight and redo the topology split. The new split will *either* assign devices with performance-managed network adapters more evenly across the domains *or* include less devices with managed or unmanaged network adapters in the domain that has a lot of devices with performance-managed network adapters. Either way, the workload across the domains will be more balanced.
ErrorRange	Floating point number Default: 0.05	Used to calculate the desired total weight for a domain. The desired maximum, (1.0 + ErrorRange) *times* (average Weight), is usually calculated by summing the number of network adapters *unless* that number is modified by ManagedAdapterWeight, UnmanagedAdapterWeight, and PerfAdapterWeight.

Table 9-2. Attributes for TopoSplit_Manager::TopoSplit-Manager (Continued)

Attribute	Value	Description
BalanceThreshold	Floating point number Default: 2.0	Determines whether the balance of the topology split across the domains is acceptable (ImbalanceRatio *less than or equal to* BalanceThreshold) or unacceptable (ImbalanceRatio *greater than* BalanceThreshold). If unacceptable and the number of retries (incremental splits) as specified in the NumberOfRetries attribute is reached without an acceptable balanced split, the event “Imbalance” is generated and written to the Topology Split Manager’s log file. The message that accompanies the event suggests that the user should increase the BalanceThreshold or the NumberOfRetries. **Note** ImbalanceRatio is a computed attribute and therefore is not user-configurable. Its value is the number of devices in the maximum-sized domain *divided by* the number of devices in the minimum-sized domain. In trial learning mode, BalanceThreshold applies to each trial topology split. If the balance of a trial topology split is unacceptable after NumberOfRetries, the trial is discarded.
NumberOfRetries	Integer Default: 3	Maximum number of attempts to split the topology such that the ImbalanceRatio of the split does not exceed the BalanceThreshold. If after “NumberOfRetries” attempts the ImbalanceRatio is still above the BalanceThreshold, the split is considered unsuccessful.

In this release, the InitiateDiscovery attribute for TopoSplit_Manager is no longer supported. As a result, the TopoSplit_Manager can no longer can be configured to initiate IP Availability Manager pending discovery at the end of a topology synchronization.

The following sections clarify the use of the parameters in the tsm.conf file:

- *Enabling topology split on page 65*
- *Enabling trial learning mode on page 66*
- *Enabling automatic topology synchronization on page 67*
- *Enabling edge management on page 68*
- *Enabling the TSM hook script on page 69*

Description of topo-split.conf

The topo-split.conf file contains configuration parameters that you edit to specify domain names for topology split. The parameters are a subset of table entries that are defined for the TopoSplit_Manager object.

When topology split is invoked, the Topology Split Manager loads the topo-split.conf file.

The topo-split.conf file contains the following active parameters:

```
DOMAINNAME=AM001
DOMAINNAME=AM002
```

It also contains the following commented-out parameters:

```
# SEED=<IPorName> DOMAIN=<DomainName>
# PATTERN=<Pattern> DOMAIN=<DomainName>
# DOMAINNAME=<DomainName>
```

The parameters and their values are case-sensitive.

In this release, only the DOMAINNAME parameter is supported. That is, you can no longer use the SEED or PATTERN parameters to assign devices to domains.

You create an entry, DOMAINNAME=<DomainName>, for each domain to which topology will be split. *"Specifying domain names for topology split" on page 63* clarifies the use of the DOMAINNAME parameter.

Description of override-polling-system.txt

You use the override-polling-system.txt partition.conf file to specify devices that the IP Availability Managers should poll at a one-minute interval. Each entry that is created includes the hostname and IP address of the device that is to be polled at a one-minute interval.

When topology split is invoked and the NumberOfDomains and IsTSMHookEnabled configuration parameters are set, the Topology Split Manager loads the override-polling-system.txt file.

The override-polling-system.txt file contains the following commented-out entry line:

```
# <HOST> <IP>
```

You create an entry, <HOST> <IP>, for each device that is eligible for monitoring at a one-minute interval, where <HOST> is hostname of the device, and <IP> is the IP address of the device. *"Specifying a list of devices for one-minute polling" on page 64* clarifies the use of the <HOST> <IP> entry line.

Tuning Topology Split Manager discovery

After copying the discovery configurations from one of the IP Availability Manager installations to the Topology Split Manager installation, you should consider modifying certain configurations in the discovery.conf and tpmgr-param.conf files to reduce memory usage and to shorten the discovery time of the Topology Split Manager. If neither high memory usage nor long discovery time is an issue in your deployment, you can ignore the following considerations.

Consider doing the following to reduce memory usage and to shorten discovery time:

- Increase the number of discovery threads in discovery.conf *Increase the number of discovery threads in discovery.conf*

- Disable unnecessary discovery probes in discovery.conf *Disable unnecessary discovery probes in discovery.conf*
- Ensure that unnecessary discovery probes are disabled in tpmgr-param.conf *Ensure that unnecessary discovery probes are disabled in tpmgr-param.conf*

  Also, consider adding the **--**nodx option to the Topology Split Manager service command, to reduce Topology Split Manager memory use by preventing the loading of the libraries that enable codebook correlation and event generation. If necessary, use property subscription to subscribe to DiscoveryError events for UnitaryComputerSystem elements.

  If memory usage is not an issue, do not add the **--**nodx option to the Topology Split Manager service command.

Increase the number of discovery threads in discovery.conf

If you are experiencing slow response SNMP discovery polling and your current CPU usage is relatively low (5 to 10 percent), consider configuring the following parameter in the discovery.conf file to a larger number to improve your discovery performance:

numberProbeThreads (Default: 10)

Increasing the number of discovery threads to 20, 30, or even 50 is acceptable, but remember that more threads might require additional or more capable CPUs, and that the requirements during discovery will increase. Increase discovery threads conservatively while monitoring SNMP Processing Statistics changes, as described in the VMware Smart Assurance IP Manager Deployment Guide.

The assumption here is that network latency is high and CPU usage is low. if CPU usage is already high during discovery, increasing the number of discovery threads might not improve discovery performance, and might even decrease discovery performance because of thread contention.

To increase the number of discovery threads for the Topology Split Manager:

1. Go to the TSM_BASEDIR/smarts/bin directory and enter the following command:

   ```
   sm_edit conf/discovery/discovery.conf
   ```

2. Locate the numberProbeThreads parameter and change its value to a larger value.

   For example:

   ```
   numberProbeThreads = 20
   ```

3. Save and close the file.

   The modified version of the discovery.conf file is saved to the TSM_BASEDIR/smarts/local/conf/discovery directory.

4 If the Topology Split Manager was running before you edited the discovery.conf file, either restart the Topology Split Manager or invoke the following dmctl command:

```
sm_tpmgr -s <TSM name> --load-conf=discovery.conf
```

Disable unnecessary discovery probes in discovery.conf

To improve the performance of the Topology Split Manager, consider configuring the following parameters in the discovery.conf file to FALSE to disable their associated discovery probes:

- STPTrunkEnabled (Default: TRUE) (not currently in discovery.conf)
- HealthProbeEnabled (Default: TRUE) (not currently in discovery.conf)
- IPv6DiscoveryEnabled (Default: TRUE)
- HostResourcesEnabled (Default: TRUE)
- VirtualRouterEnabled (Default: TRUE)
- EnableF5PoolDiscovery (Default: TRUE) (not currently in discovery.conf)

  The STP probe, which is associated with STPTrunkEnabled, discovers Layer 2 connectivity by using spanning tree protocol information. If only a few connections in your network use STP, you should disable the STP probe to reduce memory consumption and discovery time.

  Also, consider configuring the following parameters in the discovery.conf file to TRUE to reduce the amount of SNMP information that is retrieved for network adapters and to enable SNMP get-bulk requests:

- MinimumDemandGet (Default: FALSE)
- SNMPGetBulkEnabled (Default: TRUE)

  To edit the discovery.conf file:

  a Go to the TSM_BASEDIR/smarts/bin directory and enter the following command:

  ```
  sm_edit conf/discovery/discovery.conf
  ```

  b Locate the parameters that you want to change and change their values.

  For example:

  ```
  IPv6DiscoveryEnabled = FALSE
  ```

  c Save and close the file.

  The modified version of the discovery.conf file is saved to the TSM_BASEDIR/smarts/local/conf/discovery directory.

d If the Topology Split Manager was running before you edited the discovery.conf file, either restart the Topology Split Manager or invoke the following dmctl command:

```
sm_tpmgr -s <TSM name> --load-conf=discovery.conf
```

Ensure that unnecessary discovery probes are disabled in tpmgr-param.conf

To improve the performance of the Topology Split Manager, ensure that the following parameters in the tpmgr-param.conf file are set to TRUE to disable their associated discovery probes:

- DisableVirtualTCPDiscovery (Default: TRUE)
- DisableMPLSDiscovery (Default: TRUE)
- DisableMulticastDiscovery (Default: TRUE)
- DisableProtocolDiscovery (Default: TRUE)
- DisableVOIPDiscovery (Default: TRUE)

  The VirtualTCP probe is only necessary for deployments that include the VMware Smart Assurance Application Connectivity Monitor. Objects that are discovered by the MPLS, Multicast, Network Protocol, and VoIP probes are not necessary for the topology split.

  To edit the tpmgr-param.conf file:

  a Go to the TSM_BASEDIR/smarts/bin directory and enter the following command:

  ```
  sm_edit conf/discovery/tpmgr-param.conf
  ```

  b Locate the parameters that you want to change and change their values.

  For example:

  ```
  DisableMPLSDiscovery TRUE
  ```

  c Save and close the file.

  The modified version of the tpmgr-param.conf file is saved to the TSM_BASEDIR/smarts/local/conf/discovery directory.

  d If the Topology Split Manager was running before you edited the tpmgr-param.conf file, either restart the Topology Split Manager or invoke the following dmctl command:

  ```
  sm_tpmgr -s <TSM name> --load-conf=tpmgr-param.conf
  ```

Specifying domain names for topology split

By default, the Topology Split Manager is configured to split the topology into two domains named AM001 and AM002. The name of a domain and the name of the IP Availability Manager that is managing the domain are the same name.

You can edit the topo-split.conf file to change these names or to configure additional IP Availability Managers as destinations for the Topology Split Manager. The parameter definitions to be edited in the topo-split.conf file are:

```
DOMAINNAME=AM001
DOMAINNAME=AM002
```

The name of each IP Availability Manager to which the topology will be split must be specified in the topo-split.conf file.

To change the default IP Availability Manager names, or to add additional IP Availability Managers as destinations for the Topology Split Manager:

1 Go to the TSM_BASEDIR/smarts/bin directory and enter the following command to open the topo-split.conf file:

```
sm_edit conf/discovery/topo-split/topo-split.conf
```

2 Find the following parameter definitions at the end of the file:

```
DOMAINNAME=AM001
DOMAINNAME=AM002
```

3 Change AM001 to the name of your first IP Availability Manager destination, and change AM002 to the name of your second IP Availability Manager destination.

4 To specify additional IP Availability Manager destinations, copy and paste a DOMAINNAME=<DomainName> line for each IP Availability Manager, and then change <DomainName> to the name of the IP Availability Manager.

5 Save and close the file.

   The modified version of the topo-split.conf file is saved to the BASEDIR/smarts/local/conf/discovery/topo-split directory.

   When topology split is invoked, the Topology Split Manager loads the topo-split.conf file.

Specifying a list of devices for one-minute polling

By default, the Topology Split Manager does not send the IP Availability Managers a list of devices that are eligible for monitoring at a one-minute interval.

For the Topology Split Manager to learn which devices are eligible for polling at a one-minute interval, you must add them to the override-polling-system.txt file.

The syntax for a device entry is:

```
<HOST> <IP>
```

For example:

```
kirk.vmware.com 10.10.9.9
```

To add device entries to the override-polling-system.txt:

1 Go to the TSM_BASEDIR/smarts/bin directory and enter the following command to open the override-polling-system.txt file:

```
sm_edit conf/custom/override-polling-system.txt
```

2 Add your device entries, one entry per line.

3 Save and close the file.

The modified version of the override-polling-system.txt file is saved to the TSM_BASEDIR/smarts/local/conf/custom directory.

When topology split is invoked, and assuming that topology split is enabled and the TSM hook script is enabled, as explained in *"Enabling topology split" on page 65* and *"Enabling the TSM hook script" on page 69*, the Topology Split Manager loads the override-polling-system.txt file and sends the IP Availability Managers a list of devices that should be polled at a one-minute interval.

Enabling topology split

By default, topology split is disabled.

You can edit the tsm.conf file and enable this feature. The line to be edited is:

```
NumberOfDomains = 1
```

To enable topology split:

1 Go to the TSM_BASEDIR/smarts/bin directory and enter the following command to open the tsm.conf file:

```
sm_edit conf/discovery/topo-split/tsm.conf
```

2 Find the following line:

```
NumberOfDomains = 1
```

3 Change “1” to the number of domains into which the topology will be split. That number must be equal to the number of domains that are specified in the TSM_BASEDIR/smarts/local/conf/discovery/topo-split/topo-split.conf file.

4 Save and close the file.

The modified version of the tsm.conf file is saved to the TSM_BASEDIR/smarts/local/conf/discovery/topo-split directory.

Or, as an alternative to the preceding procedure, and assuming that the NumberOfDomains parameter is commented-out in the tsm.conf file, use the dmctl utility to set the parameter. Go to the TSM_BASEDIR/smarts/bin directory and issue the following command:

```
dmctl -s
<TSM name>
put TopoSplit_Manager::TopoSplit-Manager::NumberOfDomains

<number of domains into which to split the topology>
```

Enabling trial learning mode

By default, trial learning mode is disabled.

You can edit the tsm.conf file and enable this feature. The lines to be edited are:

```
NumberOfDomains = 1
# IsTrialLearningEnabled = FALSE
# IsSplitFromScratch = FALSE
# NumberOfTrials = 3
```

To enable trial learning mode:

1 Go to the TSM_BASEDIR/smarts/bin directory and enter the following command to open the tsm.conf file:

```
sm_edit conf/discovery/topo-split/tsm.conf
```

2 Find the following lines:

```
NumberOfDomains = 1
# IsTrialLearningEnabled = FALSE
# IsSplitFromScratch = FALSE
# NumberOfTrials = 3
```

3 Make the following edits to the four parameters:

```
NumberOfDomains = <number of domains into which to split the topology>
IsTrialLearningEnabled = TRUE
IsSplitFromScratch = TRUE
NumberOfTrials = <number of initial (full) topology splits to run>
```

4 Save and close the file.

The modified version of the tsm.conf file is saved to the TSM_BASEDIR/smarts/local/conf/discovery/topo-split directory.

Or, as an alternative to the preceding procedure, and assuming that the four parameters are commented-out in the tsm.conf file, use the dmctl utility to set these parameters. Go to the TSM_BASEDIR/smarts/bin directory and issue the following commands:

```
            dmctl -s
            <TSM name>
            put TopoSplit_Manager::TopoSplit-Manager::NumberOfDomains

<number of domains into which to split the topology>
            dmctl -s
            <TSM name>
            put TopoSplit_Manager::TopoSplit-Manager::IsTrialLearningEnabled TRUE
            dmctl -s
            <TSM name>
            put TopoSplit_Manager::TopoSplit-Manager::IsSplitFromScratch TRUE
            dmctl -s
            <TSM name>
            put TopoSplit_Manager::TopoSplit-Manager::NumberOfTrials

<number of initial (full) topology splits to run>
```

Enabling automatic topology synchronization

By default, automatic topology synchronization is disabled.

You can edit the tsm.conf file and enable this feature. The lines to be edited are:

```
NumberOfDomains = 1
# IsTopoSyncEnabled = TRUE
# IsTrialLearningEnabled = FALSE
```

To enable automatic topology synchronization:

1 Go to the TSM_BASEDIR/smarts/bin directory and enter the following command to open the tsm.conf file:

```
            sm_edit conf/discovery/topo-split/tsm.conf
```

2 Find the following lines:

```
NumberOfDomains = 1
# IsTopoSyncEnabled = TRUE
# IsTrialLearningEnabled = FALSE
```

3 Make the following edits to the three parameters:

```
NumberOfDomains = <number of domains into which to split the topology>
IsTopoSyncEnabled = TRUE
IsTrialLearningEnabled = FALSE
```

4 Save and close the file.

The modified version of the tsm.conf file is saved to the TSM_BASEDIR/smarts/local/conf/discovery/topo-split directory.

Or, as an alternative to the preceding procedure, and assuming that the three parameters are commented-out in the tsm.conf file, use the dmctl utility to set these parameters. Go to the TSM_BASEDIR/smarts/bin directory and issue the following commands:

```
              dmctl -s
              <TSM name>
              put TopoSplit_Manager::TopoSplit-Manager::NumberOfDomains

   <number of domains into which to split the topology>
              dmctl -s
              <TSM name>
              put TopoSplit_Manager::TopoSplit-Manager::IsTopoSyncEnabled TRUE
              dmctl -s
              <TSM name>
              put TopoSplit_Manager::TopoSplit-Manager::IsTrialLearningEnabled FALSE
```

Enabling edge management

By default, edge management is disabled.

You can edit the tsm.conf file and enable this feature. The lines to be edited are:

```
NumberOfDomains = 1
# EnableAutoEdgeManagement = TRUE
# IsTopoSyncEnabled = TRUE
# IsTrialLearningEnabled = FALSE
```

To enable edge management:

1 Go to the TSM_BASEDIR/smarts/bin directory and enter the following command to open the tsm.conf file:

```
sm_edit conf/discovery/topo-split/tsm.conf
```

2 Find the following lines:

```
NumberOfDomains = 1
# EnableAutoEdgeManagement = TRUE
# IsTopoSyncEnabled = TRUE
# IsTrialLearningEnabled = FALSE
```

3 Make the following edits to the four parameters:

```
NumberOfDomains = <number of domains into which to split the topology>
EnableAutoEdgeManagement = TRUE
IsTopoSyncEnabled = TRUE
IsTrialLearningEnabled = FALSE
```

4 Save and close the file.

The modified version of the tsm.conf file is saved to the TSM_BASEDIR/smarts/local/conf/discovery/topo-split directory.

Or, as an alternative to the preceding procedure, and assuming that the four parameters are commented-out in the tsm.conf file, use the dmctl utility to set these parameters. Go to the TSM_BASEDIR/smarts/bin directory and issue the following commands:

```
dmctl -s
<TSM name>
put TopoSplit_Manager::TopoSplit-Manager::NumberOfDomains

<number of domains into which to split the topology>
dmctl -s
<TSM name>
put TopoSplit_Manager::TopoSplit-Manager::EnableAutoEdgeManagement TRUE
dmctl -s
<TSM name>
put TopoSplit_Manager::TopoSplit-Manager::IsTopoSyncEnabled TRUE
dmctl -s
<TSM name>
put TopoSplit_Manager::TopoSplit-Manager::IsTrialLearningEnabled FALSE
```

Enabling the TSM hook script

By default, the TSM hook script is disabled.

You can edit the tsm.conf file and enable this feature. The lines to be edited are:

```
NumberOfDomains = 1
IsTSMHookEnabled = FALSE
# IsTrialLearningEnabled = FALSE
```

To enable the TSM hook script:

1 Go to the TSM_BASEDIR/smarts/bin directory and enter the following command to open the tsm.conf file:

```
sm_edit conf/discovery/topo-split/tsm.conf
```

2 Find the following lines:

```
NumberOfDomains = 1
IsTSMHookEnabled = FALSE
# IsTrialLearningEnabled = FALSE
```

3 Make the following edits to the three parameters:

```
NumberOfDomains = <number of domains into which to split the topology>
IsTSMHookEnabled = TRUE
IsTrialLearningEnabled = FALSE
```

4 Save and close the file.

The modified version of the tsm.conf file is saved to the TSM_BASEDIR/smarts/local/conf/discovery/topo-split directory.

Or, as an alternative to the preceding procedure, and assuming that the three parameters are commented-out in the tsm.conf file, use the dmctl utility to set these parameters. Go to the TSM_BASEDIR/smarts/bin directory and issue the following commands:

```
dmctl -s
<TSM name>
put TopoSplit_Manager::TopoSplit-Manager::NumberOfDomains

<number of domains into which to split the topology>
dmctl -s
<TSM name>
put TopoSplit_Manager::TopoSplit-Manager::IsTSMHookEnabled TRUE
dmctl -s
<TSM name>
put TopoSplit_Manager::TopoSplit-Manager::IsTrialLearningEnabled FALSE
```

Configuring security

The security configuration files, clientConnect.conf and serverConnect.conf, located in the BASEDIR/smarts/conf directory of each VMware Smart Assurance software installation, enable you to set up secure connections between the component applications in a Topology Split Manager deployment. By default, the configuration option settings in the clientConnect.conf file and the serverConnect.conf files enable minimally secure connections between the components.

The VMware Smart Assurance System Administration Guide presents detailed information about VMware Smart Assurance secure communications.

Initiating Discovery and Topology Split

10

This chapter includes the following topics:

- Overview
- Topology Split Manager best practices
- Preparing for discovery
- Starting the Broker and the Managers
- Performing an initial full discovery in trial learning mode
- Performing an initial full discovery in normal mode
- Performing an incremental discovery
- Initiating a discovery
- Initiating a full topology split
- Initiating an incremental topology split
- Using an initial topology split from trial learning mode

Overview

A full discovery with full topology split should be used only for an initial discovery or after a major network redesign. Thereafter, incremental discovery with incremental topology split should be used for normal topology updates.

An initial full discovery with full topology split can take a long time to complete. It consists of the time that the Topology Split Manager takes to discover the topology in the managed IP network *plus* the time that the topology-split algorithm takes to assign the discovered devices to the IP Availability Managers *plus* the time that the IP Availability Managers take to discovery the devices that have been assigned to them. While discoveries on the IP Availability Managers are in progress, the IP Availability Managers cannot monitor the network.

Topology Split Manager best practices

Full discovery with full topology split is done initially and then on an as-needed basis, typically after a major network redesign. Incremental discovery with incremental topology split is done on a regular basis, which could be daily, twice a week, or some other interval, depending on the customer's environment. The duration of an incremental discovery is much shorter than the duration of a full discovery.

Autodiscovery for the Topology Split Manager could be enabled for full discovery. Autodiscovery for the IP Availability Managers should be disabled for both full discovery and incremental discovery.

Full discovery operation and guidelines

During the initial full discovery with full topology split, the Topology Split Manager discovers the topology in the managed IP network and then assigns the discovered devices to the IP Availability Managers. During any subsequent full discovery with full topology split, the Topology Split Manager rediscovers the topology in the managed IP network and then reassigns the discovered devices to the IP Availability Managers.

Rediscovering the topology in the managed IP network reveals any changes in the configuration or connectivity of the managed devices and keeps the topology current. For example, if a card is added to a switch, the Topology Split Manager must rediscover the switch so that it can add the new card to the topology. And if new devices have been added to the network and autodiscovery is enabled, autodiscovery will discover the new devices during the rediscovery.

Incremental discovery operation and guidelines

During an incremental discovery with incremental topology split, the Topology Split Manager discovers topology updates and distributes those updates to the IP Availability Managers. Topology updates consist of any of the following:

- Devices that have been added to the managed IP network since the last Topology Split Manager discovery
- Devices that have been decommissioned (removed) from the managed IP network since the last Topology Split Manager discovery
- Configuration changes that have occurred in the managed IP network since the last Topology Split Manager discovery

  New devices can be added to the Topology Split Manager as seed entries in a seed file or placed on the Pending Devices list of the Topology Split Manager. Decommissioned devices can be removed from the Topology Split Manager through the Global Console and the **Delete** command or by calling the VMware Smart Assurance **remove()** method.

  VMware envisions that customers learn about topology updates through their automated service order provisioning systems, and use cron (or sm_sched scheduler) and custom scripts that do the following at a regular interval:

- Extract from the provisioning system a list of devices that have been added to the system since the previous query and write that list to a table, say AddList.
- Extract from the provisioning system a list of devices that have been decommissioned (removed) from the system since the previous query and write that list to a table, say RemoveList.
- Call **TopoSplit_Manager::TopoSplit-Manager::addPending()** for each device in AddList to add the devices in AddList to the Topology Split Manager's Pending Devices list.
- Call **remove()** for each device in RemoveList to remove the devices (and their components and connections) in RemoveList from the Topology Split Manager's repository.
- Invoke the following command to initiate a Topology Split Manager pending discovery:

```
dmctl -s <TSM name> invoke  ICF_TopologyManager::ICF-TopologyManager

   discoverPending
```

In response, the Topology Split Manager performs a pending discovery to discover the devices on its Pending Devices list, which triggers an incremental topology split. During the incremental topology split, the Topology Split Manager performs the following actions:

- Assigns the new devices to the appropriate IP Availability Managers and calls **ICF_TopologyManager::ICF-TopologyManager::addPending()** to add the new devices to the appropriate IP Availability Managers' Pending Devices list.
- Calls **remove()** for each decommissioned device to remove the decommissioned devices (and their components and connections) from the topologies of the appropriate IP Availability Managers.

Preparing for discovery

Preparing a Topology Split Manager deployment for discovery requires configuring all component applications in the deployment so that they can form the proper physical and logical connections to one another. Some of these tasks are performed before the components are started, such as the editing of configuration files, while others are performed through the Global Console when the components are up and running.

Prepare the Topology Split Manager for discovery

Ensure that the Topology Split Manager is configured as the IP Availability Managers are configured and that certain discovery configurations have been fine tuned to improve Topology Split Manager discovery.

The procedures for configuring discovery are given in the
VMware Smart Assurance IP Manager User Guide and the
VMware Smart Assurance IP Manager Reference Guide. The procedures for fine tuning discovery and for configuring topology split are given in Chapter 9 Configuring the Topology Split Manager

You, as an administrator, might want to enable autodiscovery in the Topology Split Manager for full discovery, but not for incremental discovery. For incremental discovery, the Topology Split Manager should discover and process just the topology updates that you present to the Topology Split Manager.

If you have access to a topology source, such as a hardware-inventory database, from which you a can easily create a seed file that contains a complete list of the network topology that is to be discovered, you should consider manual discovery *without* autodiscovery. When using manual discovery without autodiscovery, the Topology Split Manager discovers just the seed systems (seed devices) that are specified in the seed file.

If the network topology information is incomplete, unavailable, or constantly changing, you should consider manual discovery *with* autodiscovery. The seed file can be much smaller because the seed systems that are specified for manual discovery also serve as a starting point for autodiscovery. Autodiscovery automatically discovers the managed network from the seed systems.

Note You create discovery filters and (optional) specify autodiscovery safeguards to ensure that autodiscovery does not automatically add to the modeled topology those devices that you do not intend to manage.

For manual discovery *without* autodiscovery or manual discovery *with* autodiscovery, you can use the template seed file named *seedfile*, in the TSM_BASEDIR/smarts/conf directory, as the basis for creating a seed file. By default, the Topology Split Manager expects to find the seed file in the TSM_BASEDIR/smarts/local/conf directory.

The VMware Smart Assurance IP Manager User Guide contains the procedures for creating seed files and for configuring and enabling autodiscovery.

Prepare the IP Availability Manager s for discovery

Ensure that the IP Availability Managers are configured identically and that autodiscovery is disabled. Because autodiscovery is disabled, each IP Availability Manager will discover only the devices on its Pending Devices list.

The procedures for configuring discovery are given in the VMware Smart Assurance IP Manager User Guide and the VMware Smart Assurance IP Manager Reference Guide. Also see the configuration procedures in Chapter 9 Configuring the Topology Split Manager

Prepare the Global Manager for discovery

Ensure that the Global Manager is configured to import topology and events from the IP Availability Managers.

The procedures for configuring the Global Manager with the proper DXA files are given in the VMware Smart Assurance Service Assurance Manager Configuration Guide. Also see the configuration procedures in Chapter 9 Configuring the Topology Split Manager

Prepare the optional Adapter for Alcatel-Lucent 5620 SAM EMS for discovery

Ensure that the optional Adapter for Alcatel-Lucent 5620 SAM EMS (Adapter) is configured to export Alcatel-Lucent topology to the Topology Split Manager: Set the AMServerName parameter value in the Adapter's emsConfig.import file to the name of the Topology Split Manager.

The procedures for configuring the Adapter to export Alcatel-Lucent topology to the Topology Split Manager are given in the VMware Smart Assurance Adapter for Alcatel-Lucent 5620 SAM EMS User Guide. Also see the configuration procedures in Chapter 9 Configuring the Topology Split Manager

Starting the Broker and the Manager s

To start the Broker and the Managers in the Topology Split Manager deployment:

1 Start the Broker, the Global Manager, and the IP Availability Managers as described in the VMware Smart Assurance Installation Guide for SAM, IP, ESM, MPLS, and NPM Managers.

2 Start the Topology Split Manager as described in Service command.

3 Go to the TSM_BASEDIR/smarts/bin directory and enter the following command to verify that the Managers are running:

```
brcontrol
```

Performing an initial full discovery in trial learning mode

To perform an initial full discovery with full topology split in trial learning mode:

1 Configure the Topology Split Manager as follows:

- Enable topology split

  Set the NumberOfDomains configuration parameter to the number of domains that are specified in the TSM_BASEDIR/smarts/local/conf/discovery/topo-split/ topo-split.conf file. That number must be greater than 1.

  - Select trial learning mode

    Set IsTrialLearningEnabled to TRUE and set the NumberOfTrials configuration parameter to the number of initial (full) topology splits to run. That number must be greater than 1.

  - Select full topology split

    Set the IsSplitFromScratch configuration parameter to TRUE.

2 Create a seed file to seed the Topology Split Manager topology.

   The Topology Split Manager uses the hostnames or IP addresses of the seed devices that are specified in the seed file to discover the devices in the managed IP network. The seed file resides in the TSM_BASEDIR/smarts/local/conf directory. The VMware Smart Assurance IP Manager User Guide includes the procedure for creating seed files.

3 Import the seed file and initiate a full discovery with full topology split.

   Initiating a discovery describes various methods to import a seed file and initiate a discovery. One method is to go to the TSM_BASEDIR/smarts/bin directory and enter the following command:

```
dmctl -s <TSM name> invoke

    ICF_TopologyManager::ICF-TopologyManager importFromFile

      <seed file name>
```

4 Use the resulting seed files, csv reports, and trial-hook-created reports (if any) to choose an initial topology split. Here are some guidelines for evaluating and choosing the best topology split:

   - A split with a lower imbalance ratio is preferable to a split with a higher imbalance ratio.
     - A split with a lower total proxy percentage is preferable to one with a higher total proxy percentage.

       The true load that a topology puts on a domain depends on other statistics as well.

5 If the results do not fall into your predefined parameters of performance tolerances, set IsSplitFromScratch to TRUE. Then go to the TSM_BASEDIR/smarts/bin directory and enter the following command to create new trial initial topology splits:

```
dmctl -s <TSM name> invoke GA_Driver::TopoSplit-Driver start
```

   After choosing a trial initial topology split, perform the procedure in Using an initial topology split from trial learning mode to continue with the desired trial result as a starting point.

Performing an initial full discovery in normal mode

Typically, you will perform an initial full discovery in trial learning mode to create a number of initial full topology splits, then compare the initial splits and choose the one that best matches your criteria. But if for some reason you want to create just one initial full topology split, you can use the following procedure to do so.

To perform an initial full discovery with full topology split in normal mode:

1 Configure the Topology Split Manager as follows:

- Enable topology split

  Set the NumberOfDomains configuration parameter to the number of domains that are specified in the TSM_BASEDIR/smarts/local/conf/discovery/topo-split/ topo-split.conf file. That number must be greater than 1.

  - Select normal mode

    Ensure that the IsTrialLearningEnabled, IsTopoSyncEnabled, and IsTSMHookEnabled configuration parameters are FALSE, to operate in normal mode and to disable automatic topology synchronization and the TSM hook script.

  - Select full topology split

    Set the IsSplitFromScratch configuration parameter to TRUE.

2 Create a seed file to seed the Topology Split Manager topology.

The Topology Split Manager uses the hostnames or IP addresses of the seed devices that are specified in the seed file to discover the devices in the managed IP network. The seed file resides in the TSM_BASEDIR/smarts/local/conf directory. The VMware Smart Assurance IP Manager User Guide includes the procedure for creating seed files.

3 Import the seed file and initiate a full discovery with full topology split.

Initiating a discovery describes various methods to import a seed file and initiate a discovery. One method is to go to the TSM_BASEDIR/smarts/bin directory and enter the following command:

```
sm_tpmgr -s <TSM name> --seed=<seed file name>
```

4 Use the new seed files to verify the split quality: Count and verify statistics such as the number of devices in each domain and the number of NetworkAdapters (managed/unmanaged) in each domain, to ensure that no IP Availability Manager will be overburdened.

5 If the results do not fall into your predefined parameters of performance tolerances, set IsSplitFromScratch to TRUE. Then go to the TSM_BASEDIR/smarts/bin directory and enter the following command to run another full topology split:

```
dmctl -s <TSM name> invoke GA_Driver::TopoSplit-Driver start
```

After accepting a topology split, perform the procedure in Initiating an incremental topology split to automatically distribute the topology to the IP Availability Managers.

Performing an incremental discovery

After a full discovery with full topology split, you should use incremental discovery with incremental topology split for topology updates. Topology updates are due to the adding of new devices to the managed IP network, the decommissioning of devices, or the changing of configurations.

Incremental discovery is invoked by adding new devices as seed devices to the Topology Split Manager or placing the new devices on the Topology Split Manager's Pending Devices list. Seed devices are specified in a seed file or in **Add Agent** commands.

Incremental topology split assigns the new devices to an existing topology split, which might result in the moving of some devices from one IP Availability Manager to another in an attempt to improve the split.

To perform an incremental discovery with incremental topology split:

1 Configure the Topology Split Manager as follows:

- Enable topology split

  Set the NumberOfDomains configuration parameter to the number of domains that are specified in the TSM_BASEDIR/smarts/local/conf/discovery/topo-split/ topo-split.conf file. That number must be greater than 1.

  - Select normal mode

    Ensure that the IsTrialLearningEnabled configuration parameter is FALSE.

  - Select incremental topology split

    Ensure that the IsSplitFromScratch configuration parameter is FALSE.

  - Optional but recommended: Enable automatic topology synchronization

    Set the IsTopoSyncEnabled configuration parameter to TRUE.

  - Optional but recommended: Enable edge management

    Set the EnableAutoEdgeManagement configuration parameter to TRUE. Note that automatic topology synchronization must be enabled before edge management can be enabled.

  - Enable the TSM hook script

    Set the IsTSMHookEnabled configuration parameter to TRUE.

2 Create a seed file and add the new devices to the seed file.

  The Topology Split Manager uses the hostnames or IP addresses of the seed devices that are specified in the seed file to discover the devices in the managed IP network. The seed file resides in the TSM_BASEDIR/smarts/local/conf directory. The VMware Smart Assurance IP Manager User Guide includes the procedure for creating seed files.

3 Import the seed file and initiate an incremental discovery with incremental topology split.

Initiating a discovery describes various methods to import a seed file and initiate a discovery. One method is to open a Domain Manager Administration Console that is attached to the Topology Split Manager and do the following:

Select **Topology > Import from seed file** to launch the **Import From Seed File** dialog box. In the dialog box, specify the complete path and name of the seed file. (By default, the path to the seed file is TSM_BASEDIR/smarts/local/conf.) Click **OK** to import the seed devices and to start the incremental discovery with incremental topology split.

If IsTopoSyncEnabled is TRUE, the Topology Split Manager splits the new topology and places a seed file for each domain in the Topology Split Manager's local/logs directory, and in each IP Availability Manager's local/logs directory. The Topology Split Manager adds the new devices in a domain's seed file to the Pending Devices list of the appropriate IP Availability Manager. It also deletes outdated devices from the appropriate IP Availability Manager's repository. Outdated devices are devices that are already in the IP Availability Manager's topology *and* are either not assigned by the topology-split algorithm or are assigned and are *unmanaged*.

If IsTopoSyncEnabled is FALSE, the Topology Split Manager splits the new topology and places a seed file for each domain in the Topology Split Manager's local/logs directory, but *not* in each IP Availability Manager's local/logs directory. Because the seed files are retained by the Topology Split Manager, administrators can use them later to initiate IP Availability Manager discovery manually.

Note If you use these seed files to initiate IP Availability Manager discovery manually, ensure that you use the set of seed files with the most recent timestamp.

Each new seed file in the Topology Split Manager's local/logs directory contains previously discovered devices and newly discovered devices, whereas each new seed file in a domain's local/logs directory contains the new devices only. Outdated devices are commented out in the domain's seed file.

4 If IsTopoSyncEnabled is TRUE: For each IP Availability Manager, go to the BASEDIR/smarts/bin directory in the appropriate IP Manager installation area and enter the following command to invoke a pending discovery:

```
dmctl -s <AM name>   invoke ICF_TopologyManager::ICF-TopologyManager discoverPending
dmctl -s AM001 invoke

  ICF_TopologyManager::ICF-TopologyManager discoverPending
```

5 If IsTopoSyncEnabled is FALSE: For each IP Availability Manager, use the associated seed file in the Topology Split Manager's local/logs directory to manually initiate a discovery at the IP Availability Manager. If the IP Availability Manager is on a different host, copy the associated seed file to that host and place it in the BASEDIR/smarts/local/conf directory.

6 After the IP Availability Managers have completed their discoveries, check all critical numbers to determine if they meet your predefined criteria.

Initiating a discovery

The VMware Smart Assurance IP Manager Concepts Guideand the VMware Smart Assurance IP Manager User Guidedescribe the various discovery methods for an IP Manager. Because the Topology Split Manageris a specially configured instance of an IP Manager, all of the discovery methods are available to the Topology Split Manager.

You can use any of the following methods to initiate a discovery in the Topology Split Manager. Be aware that if the NumberOfDomains configuration parameter is set, any of these discovery methods will trigger a topology split in the Topology Split Manager.

- From a Domain Manager Administration Console that is attached to the Topology Split Manager:
  - Select **Topology >Import from seed file**, or click the **Import from seed**file toolbar button ( ), to import a seed file that lists the seed devices that are to be added to the Topology Split Manager's topology.

    This method is used to add a large number of devices. It is typically used to perform a full discovery, but can also be used to perform an incremental discovery if a large number of new devices have been added to the managed IP network.
  - Select **Topology >Ad d A gent**, or click the **Add Agent**toolbar button ( ), to add seed devices to the Topology Split Manager's topology one at a time.

    This method is used to add a small number of devices. It is typically used to perform an incremental discovery.
  - Select **Topology >Discover All**to rediscover all managed devices in the Topology Split Manager's repository, and to discover the devices on the Topology Split Manager's Pending Devices list.
  - Select **Topology >Discover Pending**to discover just the devices on the Topology Split Manager's Pending Devices list.
  - Right-click a device on the Topology Split Manager's Pending Devices list and then select **Rediscover**from the pop-up menu to rediscover that device.
  - Click a device in the Topology Split Manager's topology tree and then select **Topology >Rediscover** to rediscover that device. You can also rediscover a device by right-clicking it and selecting **Rediscover**from the pop-up menu.

- From the command line:
  - Go to the TSM_BASEDIR/smarts/bin directory and enter the following command to start the Topology Split Manager, import a seed file that lists the seed devices that are to be added to the topology, and initiate a discovery:

```
sm_server --name=<
        TSM name>
        --config=icf
  --bootstrap=bootstrap-tsm.conf --seedfile=<
        seed file name>
              --ignore-restore-errors --noreconfigure --output --daemon
dmctl -s <
        TSM name>
        invoke
   ICF_TopologyManager::ICF-TopologyManager importFromFile
     <
        seed file name>
        sm_tpmgr -s <
        TSM name>
        --seed=<
        seed file name>
```

  - Go to the TSM_BASEDIR/smarts/bin directory and enter the following command to rediscover all managed devices in the Topology Split Manager's repository, and to discover the devices on the Topology Split Manager's Pending Devices list:

```
dmctl -s<
        TSM name>
        invoke
  ICF_TopologyManager::ICF-TopologyManager discoverAll
```

- Go to the TSM_BASEDIR/smarts/bin directory and enter the following command to discover just the devices on the Topology Split Manager's Pending Devices list:

```
dmctl -s<
        TSM name>
        invoke
  ICF_TopologyManager::ICF-TopologyManager discoverPending
```

Initiating a full topology split

To initiate a full topology split:

1 Configure the Topology Split Manager as follows:

- Enable topology split

  Set the NumberOfDomains configuration parameter to the number of domains that are specified in the TSM_BASEDIR/smarts/local/conf/discovery/topo-split/ topo-split.conf file. That number must be greater than 1.

  - Select trial learning mode or normal mode

    For trial learning mode, set IsTrialLearningEnabled to TRUE and set the NumberOfTrials configuration parameter to the number of initial (full) topology splits to run. That number must be greater than 1.

    For normal mode, ensure that the IsTrialLearningEnabled, IsTopoSyncEnabled, and IsTSMHookEnabled configuration parameters are FALSE, to operate in normal mode and to disable automatic topology synchronization and the TSM hook script.

2 Enter the following command from TSM_BASEDIR/smarts/bin to invoke a full topology split:

```
dmctl -s <TSM name> invoke GA_Driver::TopoSplit-Driver start
```

3 For trial learning mode, use the resulting seed files, csv reports, and trial-hook-created reports (if any) to choose an initial topology split. Here are some guidelines for evaluating and choosing the best topology split:

- A split with a lower imbalance ratio is preferable to a split with a higher imbalance ratio.
  - A split with a lower total proxy percentage is preferable to one with a higher total proxy percentage.

    The true load that a topology puts on a domain depends on other statistics as well.

4 For normal mode, use the new seed files to verify the split quality: Count and verify statistics such as the number of devices in each domain and the number of NetworkAdapters (managed/unmanaged) in each domain, to ensure that no IP Availability Manager will be overburdened.

5 If the results do not fall into your predefined parameters of performance tolerances, set IsSplitFromScratch to TRUE. Then go to the TSM_BASEDIR/smarts/bin directory and enter the following command to run another topology split:

```
dmctl -s <TSM name> invoke GA_Driver::TopoSplit-Driver start
```

Initiating an incremental topology split

To initiate an incremental topology split:

1 Configure the Topology Split Manager as follows:

- Enable topology split

  Set the NumberOfDomains configuration parameter to the number of domains that are specified in the TSM_BASEDIR/smarts/local/conf/discovery/topo-split/ topo-split.conf file. That number must be greater than 1.

  - Select normal mode

    Ensure that the IsTrialLearningEnabled configuration parameter is FALSE.

  - Select incremental topology split

    Ensure that the IsSplitFromScratch configuration parameter is FALSE.

  - Optional but recommended: Enable automatic topology synchronization

    Set the IsTopoSyncEnabled configuration parameter to TRUE.

  - Optional but recommended: Enable edge management

    Set the EnableAutoEdgeManagement configuration parameter to TRUE. Note that automatic topology synchronization must be enabled before edge management can be enabled.

  - Enable the TSM hook script

    Set the IsTSMHookEnabled configuration parameter to TRUE.

2 Enter the following command from TSM_BASEDIR/smarts/bin to invoke an incremental topology split:

```
dmctl -s <TSM name> invoke GA_Driver::TopoSplit-Driver start
```

3 If IsTopoSyncEnabled is TRUE: For each IP Availability Manager, go to the BASEDIR/smarts/bin directory in the appropriate IP Manager installation area and enter the following command to invoke a pending discovery:

```
dmctl -s <AM name>   invoke ICF_TopologyManager::ICF-TopologyManager discoverPending
dmctl -s AM001 invoke

  ICF_TopologyManager::ICF-TopologyManager discoverPending
```

4 If IsTopoSyncEnabled is FALSE: For each IP Availability Manager, use the associated seed file in the Topology Split Manager's local/logs directory to manually initiate a discovery at the IP Availability Manager. If the IP Availability Manager is on a different host, copy the associated seed file to that host and place it in the BASEDIR/smarts/local/conf directory.

5 After the IP Availability Managers have completed their discoveries, check all critical numbers to determine if they meet your predefined criteria.

Using an initial topology split from trial learning mode

After choosing a trial initial topology split, you should restart the Topology Split Manager with the repository file from that trial and perform an incremental topology split (with topology synchronization, edge management, and TSM hook enabled) to continue with the desired trial result as a starting point.

Here is the procedure for using a trial initial topology split. In the procedure, for simplification purposes, the Topology Split Manager is named INCHARGE-TSM, and the chosen initial topology split is located in the INCHARGE-TSM.trial.002 directory and is named INCHARGE-TSM.rps.

1 Go to the TSM_BASEDIR/smarts/bin directory and enter the following command to stop the Topology Split Manager:

For a Topology Split Manager that was started with a service command (sm_service):

```
sm_service stop ic-tsm
dmquit -s INCHARGE-TSM
```

2 Copy the repository file for the chosen initial topology split to the TSM_BASEDIR/smarts/local/repos/icf directory:

On a UNIX system:

```
cp /opt/InCharge/TSM/IP/smarts/local/logs/INCHARGE-TSM.trial.002/

  INCHARGE-TSM.rps  /opt/InCharge/IP/TSM/smarts/local/repos/icf
xcopy C:\InCharge\TSM\IP\smarts\local\logs\INCHARGE-TSM.trial.002\

  INCHARGE-TSM.rps  C:\InCharge\TSM\IP\smarts\local\repos\icf
```

3 Enter the following command from TSM_BASEDIR/smarts/bin to start the Topology Split Manager:

For a service startup command (sm_service):

```
sm_service start ic-tsm
sm_server --name=INCHARGE-TSM --config=icf

  --bootstrap=bootstrap-tsm.conf --ignore-restore-errors

    --noreconfigure --output --daemon
```

4 Configure the Topology Split Manager as follows:

- Enable topology split

  Ensure that the NumberOfDomains configuration parameter is set to the number of domains that are specified in the TSM_BASEDIR/smarts/local/conf/discovery/ topo-split/topo-split.conf file. That number must be greater than 1.

  - Select normal mode

    Ensure that the IsTrialLearningEnabled configuration parameter is FALSE.

  - Select incremental topology split

    Ensure that the IsSplitFromScratch configuration parameter is FALSE.

  - Enable automatic topology synchronization

    Set the IsTopoSyncEnabled configuration parameter to TRUE.

  - Enable edge management

    Set the EnableAutoEdgeManagement configuration parameter to TRUE. Note that automatic topology synchronization must be enabled before edge management can be enabled.

  - Enable the TSM hook script

    Set the IsTSMHookEnabled configuration parameter to TRUE.

5 Enter the following command from TSM_BASEDIR/smarts/bin to invoke an incremental topology split:

```
dmctl -s <TSM name> invoke GA_Driver::TopoSplit-Driver start
```

6 For each IP Availability Manager, go to the BASEDIR/smarts/bin directory in the appropriate IP Manager installation area and enter the following command to invoke a pending discovery:

```
dmctl -s <AM name>    invoke ICF_TopologyManager::ICF-TopologyManager
discoverPending
dmctl -s AM001 invoke

   ICF_TopologyManager::ICF-TopologyManager discoverPending
```

7 After the IP Availability Managers have completed their discoveries, check all critical numbers to determine if they meet your predefined criteria.

Topology-Split Data Model

11

This chapter includes the following topics:

- Data model
- Topology-split algorithm and data model
- Topology Split Manager events
- Topology-split objects and their attributes

Data model

The Topology Split Manager discovers and builds a data model of the devices in the managed IP network. The model represents the devices objects, their relationships, their internal components, and their logical connections.

The Topology Split Manager also builds a data model of special topology-split objects in order to split the discovered objects across multiple domains. That data model, shown in Topology-split data model, is the focus of this discussion.

Figure 11-1. Topology-split data model



The topology-split objects in the data model are represented as instances of the following VMware Smart Assurance classes:

- TopoSplit_Domain

- TopoSplit_InChargeDomain
- TopoSplit_Node
- TopoSplit_Link
- TopoSplit_SuperLink

  TopoSplit_Domain and TopoSplit_InChargeDomain are sibling classes, and TopoSplit_Node and TopoSplit_Link are sibling classes. TopoSplit_SuperLink is a subclass of TopoSplit_Link.

  The TopoSplit_Manager class, not shown in Topology-split data model, manages the TopoSplit_InChargeDomain objects. It is a singleton class (the instance is TopoSplit_Manager::TopoSplit-Manager) and has attributes for configuring various aspects of the topology split.

  The VMware Smart Assurance ICIM Reference provides descriptions of VMware Smart Assurance classes, relationships, connections, and hierarchies.

Topology-split algorithm and data model

The topology-split software uses the topology-split classes to split the discovered devices into a user-specified number of domains such that the domains are fairly equal in size and the number of proxy devices (replicated devices) is kept to a minimum. If a device is a member of two domains, it is said to be replicated *once*; if a device is a member of three domains, it is said to be replicated *twice*; and so on.

The Topology Split Manager considers the devices as elements in a set, and considers each connection between a pair of devices as a subset.

Phase 1: initial split

In phase 1, the initial split, the topology-split software assigns the subsets to the domains in a way that minimizes the number of device replications. In Phase 1: initial split, the initial device-to-domain assignments are represented by the relationships between TopoSplit_Node and TopoSplit_Domain.

Phase 2: improved split

In phase 2, a heuristic graph-partitioning algorithm named *Kernighan-Lin* attempts to improve the device-to-domain assignments by moving subsets from one domain to another, while keeping the domain sizes fairly even. The algorithm continues to move subsets until no improvement is achieved. In Phase 2: improved split, the improved and final device-to-domain assignments are represented by the relationships between TopoSplit_Node and TopoSplit_InChargeDomain.

Graph contraction

Part B in Graph contraction represents the working devices and links for *graph contraction*.

As explanation, a common abstraction for relational networks is a *graph*, in which entities are represented by *vertices*, and observed relationships between the entities are represented as *edges* that connect the vertices. In addition, the vertices and edges may be attributed by *weights* to describe the strength of the vertices and the connections.

For the managed IP network that is to be split into domains, the devices in the network are the graph's vertices, and the network connections between the devices are the graph's edges. The weight of a device is proportional to the number of connections to the device, where a higher number indicates a larger number of connections.

Other weights, or attributes, of a device are ManagedAdapterWeight, UnmanagedAdapterWeight, and PerfAdapterWeight, which pertain to a device's managed, unmanaged, and performance-managed network adapters, respectively. Performance-managed network adapters belong to devices that are to have their components polled every minute. The higher the weight for a particular type of network adapter, the higher the probability that network adapters of that type will be distributed evenly across the domains.

To construct the graph, that is, to partition the devices into domains, the Kernighan-Lin algorithm starts by selecting one device (node) for each domain. From there, the algorithm uses a nearest-neighbor approach to assign the remaining devices to the domains. The neighborhood of a device is the set of all adjacent devices, including the device, itself.

For each device, the Kernighan-Lin algorithm considers the device's neighbors and evaluates the gain in modularity that would occur by moving the device from its current domain to the domain of one of its neighbors. The device is then placed in the domain for which this gain is maximum until no individual move can improve the modularity.

For an incremental topology split, the Kernighan-Lin algorithm assigns a new device to a previously defined domain by looking only at the device's neighborhood, or the set of devices that are two hops away, without redoing a full topology split. At the same time, the algorithm tries to improve the device-to-domain mapping by moving single devices into different domains, which explains why devices are sometimes moved to different domains during an incremental topology split.

Topology Split Manager events

List of Topology Split Manager events lists the events that the Topology Split Manager is capable of generating. For each event, the table identifies the managed object for the event and the condition or threshold that will trigger the event. Because no VMware Smart Assurance application subscribes to the events, the events are not evaluated or polled and thus the state (notified, cleared) of an event never changes.

Table 11-1. List of Topology Split Manager events

Managed object	Event	Condition or threshold
TopoSplit_Manager	Imbalance	The topology-split algorithm failed to reach an acceptable balanced split (ImbalanceRatio 1 > BalanceThreshold 1) in the allowed number of attempts. The allowed number of attempts is 3 by default and is specified in NumberOfRetries 1. ImbalanceRatio is equal to MaxDomainWeight 1 divided by AvgDomainWeight 1. The topology-split algorithm stops, and the “Imbalance” event is written to the Topology Split Manager’s log file.
TopoSplit_InChargeDomain	TooManyDevices	The number of devices is greater than the DeviceHighWaterMark value (NumberOfDevices 2> DeviceHighWaterMark 2). The value of DeviceHighWaterMark is propagated from the same-named attribute of TopoSplit_Manager. The topology-split algorithm is not affected by this event, and users will not see this event because no application subscribes to it.
	TooManyAdapters	The number of managed network adapters is greater than the AdapterHighWaterMark value (NumberOfManagedNetworkAdapters 2 > AdapterHighWaterMark 2). The value of AdapterHighWaterMark is propagated from the same-named attribute of TopoSplit_Manager. The topology-split algorithm is not affected by this event, and users will not see this event because no application subscribes to it.
	TooManyProxies	The number of proxies in the domain is greater than the ProxyThreshold value (ProxyRatio 1> ProxyThreshold 1). ProxyRatio is equal to the NumberOfProxies 1 divided by the NumberOfDevices 1. The topology-split algorithm is not affected by this event, and users will not see this event because no application subscribes to it.
1**Attribute of TopoSplit_Manager, which has a single instance named TopoSplit-Manager. 2**Attribute of TopoSplit_InChargeDomain, which has an instance for each domain.		

Topology-split objects and their attributes

The sections that follow describe the following object types and identify their key attributes:

- TopoSplit_Domain *TopoSplit_Domain*
- TopoSplit_InChargeDomain *TopoSplit_InChargeDomain*
- TopoSplit_Node *TopoSplit_Node*
- TopoSplit_Link *TopoSplit_Link*
- TopoSplit_SuperLink *TopoSplit_SuperLink*

TopoSplit_Domain

A TopoSplit_Domain is a domain's IP Availability Manager. During phase 1 of the topology split, the Topology Split Manager creates a TopoSplit_Domain object for each IP Availability Manager in the Topology Split Manager deployment.

Attributes for TopoSplit_Domain

Attributes for TopoSplit_Domain lists some key attributes for TopoSplit_Domain.

Table 11-2. Attributes for TopoSplit_Domain

Attribute	Description	Allowed values
NumberOfNodes	Number of owned and proxied devices in this domain.	Integer
NumberOfProxies	Number of proxied devices in this domain.	Integer
TotalWeight	Total number of managed network adapters on owned and proxied devices in this domain.	Floating point number

TopoSplit_InChargeDomain

A TopoSplit_InChargeDomain is a domain's IP Availability Manager. During phase 2 of the topology split, the Topology Split Manager creates a TopoSplit_InChargeDomain object for each IP Availability Manager in the Topology Split Manager deployment.

Attributes for TopoSplit_InChargeDomain

Attributes for TopoSplit_InChargeDomain lists some key attributes for TopoSplit_InChargeDomain.

Table 11-3. Attributes for TopoSplit_InChargeDomain

Attribute	Description	Allowed values
AdapterHighWaterMark	Maximum number of network adapters (ports, interfaces) that should be placed into a single domain.	Integer Default: 100000
DeviceHighWaterMark	Maximum number of devices that should be placed into a single domain.	Integer Default: 5000
NumberOfDevices	Number of owned devices in this domain.	Integer
NumberOfManagedNetwork-Adapters	Number of managed network adapters in this domain.	Integer
NumberOfNodes	Number of owned and proxied devices in this domain.	Integer
NumberOfProxies	Number of proxied devices in this domain.	Integer
ProxyRatio	Maximum ratio of total proxies in all domains *to* total number of devices in the Topology Split Manager.	Floating point number Default: 0.2500
TotalPrimarilyManaged-SystemWeight	Total weight of owned, managed devices in this domain.	Floating point number

Table 11-3. Attributes for TopoSplit_InChargeDomain (Continued)

Attribute	Description	Allowed values
TotalProxyWeight	Total weight of proxied devices in this domain.	Floating point number
TotalWeight	Total number of managed network adapters on owned and proxied devices in this domain.	Floating point number

TopoSplit_Node

A TopoSplit_Node is a node (device) that represents a computer-based network entity. Important nodes are Host and the following relay devices: Hub, Bridge, Switch, and Router.

For each discovered device, the Topology Split Manager creates a TopoSplit_Node object for the device.

Attributes for TopoSplit_Node

Attributes for TopoSplit_Node lists some key attributes for TopoSplit_Node.

Table 11-4. Attributes for TopoSplit_Node

Attribute	Description	Allowed values
IsManaged	True when this node is managed.	Boolean: true or false
IsRelayDevice	True when this node is a relay device.	Boolean: true or false
NodeWeight	Weight of this node, which is the sum of weights of all managed and unmanaged network adapters for this node. Computed by the topology-split algorithm and used in computing the SuperWeight of this node.	Floating point number
SuperWeight	Weight of this node plus the sum of weights of all nodes that this node contains. Computed by the topology-split algorithm and used to help balance the topology split.	Floating point number
TopoSplit_Node_Contains_NodeWeight	Sum of weights of all nodes that this node contains. Contains/BelongsTo is a relationshipset that the topology-split algorithm uses to group nodes together. Computed by the topology-split algorithm and used in computing the SuperWeight for this node.	Floating point number

TopoSplit_Link

A TopoSplit_Link is a connection between two TopoSplit_Node objects. Each complete TopoSplit_Link connects two TopoSplit_Nodes.

Attributes for TopoSplit_Link

Attributes for TopoSplit_Link lists some key attributes for TopoSplit_Link.

Table 11-5. Attributes for TopoSplit_Link

Attribute	Description	Allowed values
IsManaged	True when this link is managed.	Boolean: true or false
SuperWeight	Sum of SuperWeights for all nodes that are connected to this link.	Floating point number

TopoSplit_SuperLink

A TopoSplit_SuperLink is a working link for graph contraction.

Attributes for TopoSplit_SuperLink

Attributes for TopoSplit_SuperLink lists a key attribute for TopoSplit_SuperLink.

Table 11-6. Attributes for TopoSplit_SuperLink

Attribute	Description	Allowed values
IsManaged	True when this super link is managed.	Boolean: true or false
SuperWeight	Sum of SuperWeights for all nodes that are connected to this super link.	Floating point number

The sm_edit Utility

12

This chapter includes the following topics:

- sm_edit
- sm_edit example

sm_edit

As part of the VMware Smart Assurance deployment and configuration process, you need to modify certain files. User modifiable files include configuration files, rule set files, templates, and seed files that contain encrypted passwords. Original versions of these files are installed into appropriate subdirectories under the BASEDIR/smarts/ directory.

The original versions of files should not be altered. If a file must be modified, a new version should be created and then stored as a local copy of the file in BASEDIR/smarts/local or one of its subdirectories.

When VMware Smart Assurance software requires one of these files, it is designed to first search for a modified file in BASEDIR/smarts/local or one of its subdirectories. If a modified version of a file is not found in the local area, VMware Smart Assurance software then searches corresponding BASEDIR/smarts directories for the original version of the file.

To ease file editing and storage, VMware, Inc. provides the sm_edit utility with every VMware Smart Assurance product. When invoked, sm_edit opens the specified file in a text editor. This utility ensures that modified files are always saved to the appropriate local area and that non-local copies of all files remain unchanged. If an appropriate subdirectory does not exist for the file you are modifying, sm_edit creates the appropriate subdirectory before saving the modified file to that location. For files with header information set for encryption, sm_edit encrypts certain fields in the file. In addition, sm_edit preserves the file permissions of modified files, which helps ensure that important configuration files are not altered by unauthorized users.

In UNIX, you can configure sm_edit to use your favorite text editor when you invoke sm_edit. The VMware Smart Assurance System Administration Guide provides instructions on how to configure the utility to use a specific editor.

sm_edit example

To use sm_edit from the command line, specify the file name and include the subdirectory under BASEDIR/smarts/local where the file resides. For example, to edit the trapd.conf, enter the following command from the BASEDIR/smarts/bin directory:

```
sm_edit conf/trapd/trapd.conf
```

In this example, sm_edit searches in the BASEDIR/smarts/local/conf/trapd directory for the trapd.conf file. If it finds the trapd.conf file, it opens the file in a text editor. If sm_edit does not find the trapd.conf file in the BASEDIR/smarts/local/conf/trapd directory, it creates a local copy of the trapd.conf file and writes it to the BASEDIR/smarts/local/conf/trapd directory.

The VMware Smart Assurance System Administration Guide provides additional information about the sm_edit utility.