EMC® Smarts® IP Management Suite

7.0

Deployment Guide

P/N 300-004-085
REV A01

EMC Corporation
Corporate Headquarters:
Hopkinton, MA 01748-9103
1-508-435-1000
www.EMC.com

Copyright © 2006 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Contents

Preface vii

Chapter 1 Overview of the EMC Smarts Deployment Process

EMC Smarts architecture 1-2
The deployment process 1-4
Phase 1: Designing the EMC Smarts deployment 1-4
Phase 2: Installing and configuring the EMC Smarts components 1-5
Phase 3: Validating the deployment 1-6
Phase 4: Tuning and maintaining the deployment 1-6
Before-you-begin checklist 1-7

Chapter 2 Gathering Information for Designing

Determine the organization's requirements 2-2
Obtain network information 2-3
Obtain network diagrams 2-3
Network priorities 2-4
Identify the types of equipment in the network 2-5
Determine requirements for installing software 2-5
Integrating existing software with EMC Smarts software 2-5
Determine number of managed network devices 2-6
Refining the estimate for sizing an EMC Smarts deployment 2-6
Determine quantities of devices for licensing 2-10
Gather network security information 2-11
What other network features affect EMC Smarts? 2-12
Architectural information checklist 2-13

Chapter 3 Designing the EMC Smarts Deployment

Documenting the deployment 3-2
The Solution Architecture Diagram 3-2
The Deployment Build Guide 3-2
Determine the required size of the EMC Smarts deployment 3-3
Advantages of larger platform equipment tiers 3-4
Partitioning networks 3-5
Multiple EMC Smarts managers on a single platform 3-5
Adding information to the Solution Architecture Diagram and Deployment Build Guide 3-6
Locating EMC Smarts managers and platforms 3-7

Security and licensing considerations 3-9
Considering volume licensing configurations 3-9
Considering security and firewalls 3-9
Considering high availability configurations 3-10
Designing for overlapping (duplicate) IP networks 3-11
Designing acceptance tests 3-12
Solution Architecture Diagram checklist 3-13

Chapter 4 Planning for Discovery

Designing discovery 4-2
Initial discovery 4-2
Topology maintenance and subsequent discovery 4-4
Discovery and security 4-6
Discovery and certified device types 4-6
Discovery and the Domain Name System 4-7
Advanced discovery post-processing 4-7
Discovery design checklist 4-8

Chapter 5 Designing Polling and Thresholds

Polling and polling groups 5-2
Considering network latency 5-2
Threshold values and threshold groups 5-3
Polling and threshold checklist 5-4

Chapter 6 Deploying Syslog Processing

Syslog processing applications 6-2
Creating the syslog file 6-2
Processing the syslog file 6-2
Syslog processing checklist (optional) 6-4

Chapter 7 Designing for Administration of EMC Smarts

Who are the Global Console users? 7-2
Users and security 7-2
Designing user profiles 7-4
Designing notification lists 7-5
Restricting console operations 7-6
Designing consoles 7-7
Planning for tools and tool deployment 7-8
Administration Design Checklist 7-9

Chapter 8 Deploying EMC Smarts

General installation/deployment guidelines 8-2
Allow access to MIBs in network devices 8-2
EMC Smarts licensing 8-3
EMC Smarts installation 8-4
Setting environment variables 8-4
Configure security 8-5
Deploy trap processing 8-6
Deploy EMC Smarts user configurations 8-7

Chapter 9 Validating Your Deployment (Acceptance Testing)

Validation techniques 9-2
Initial validation 9-3
Validating discovery 9-4
Validating polling and events 9-5
Validating trap processing 9-6
Validating users and capabilities 9-7

Chapter 10 Tuning Your Deployment to Improve Performance

Performance tuning guidelines 10-2
Assessing performance 10-3
Checking resource requirements against operating system limits 10-3
Reviewing EMC Smarts license metrics 10-4
Reviewing EMC Smarts performance metrics 10-5
Codebook tasks 10-5
Duration of last discovery 10-6
ICMP processing statistics 10-8
SNMP processing statistics 10-8
Improving performance 10-10
Other tuning issues 10-11
Adjust performance thresholds to reduce inappropriate alarms 10-11
Use batching to improve trap processing performance 10-11
Global Console performance 10-11

Appendix A Using Discovery Manager to Obtain Network Information

Deploying temporarily A-2
Deploying permanently A-3
Discovery process A-4
Exporting topology A-5
Using sm_tpmgr A-5

Appendix B Managing Overlapping IP Networks With EMC Smarts

Using IP management domains to manage overlapping IP networks B-2
Using virtual IP interfaces to direct management traffic B-2
Routing management traffic to and from the management domains B-4
Configuring devices to send SNMP traps B-5
Consolidating management domain information B-7

Appendix C Design and Deployment Checklists

Before you begin checklist C-2
Architectural information checklist C-3
Solution architecture diagram checklist C-5
Discovery design checklist C-6
Polling and threshold checklist C-8
Syslog processing checklist (optional) C-9

Index i-1

Preface

The *EMC Smarts IP Management Suite Deployment Guide* describes how to deploy the EMC Smarts IP Management Suite and related components of the EMC Smarts Service Assurance Management Suite.

This guide is intended as a comprehensive resource for deployment tasks. In many cases, you will be referred to other EMC Smarts documents for specific procedures and configuration tasks, such as:

- The *EMC Smarts IP Management Suite Installation Guide* that accompanied your software product suite for detailed hardware requirements, platform requirements, and installation procedures
- The *EMC Smarts IP Management Suite Discovery Guide* for information about discovery configuration and processing
- The *EMC Smarts IP Availability Manager User's Guide* and the *EMC Smarts IP Performance Manager User's Guide* for information about EMC Smarts Manager functions including polling and thresholds
- The *EMC Smarts Service Assurance Manager Configuration Guide* for information about the Global Manager Administration Console
- The *EMC Smarts System Administration Guide* for detailed administration procedures for the EMC Smarts deployment

Audience for this guide

This guide is intended for the following audiences:

- Network and system administrators who aid in the design and deployment of EMC Smarts and its maintenance
- Integrators and network consultants who aid in designing EMC Smarts deployments and then install, validate, and tune the deployments
- Systems Engineers who design EMC Smarts deployments
- EMC Professional Services personnel who design, install, validate, and tune EMC Smarts deployments
- EMC Support personnel who respond to inquiries, problems, and issues that arise during EMC Smarts deployments

Prerequisites

To successfully use this guide, you should be familiar the following subjects:

- TCP/IP networking concepts

- Network administration practices
- The EMC Smarts architecture
- Concepts discussed in *EMC Smarts Service Assurance Manager Introduction*

EMC Smarts IP Management Suite installation directory

In this document, the term BASEDIR represents the location where EMC Smarts software is installed.

- For UNIX, this location is: /opt/InCharge*<n>*/*<productsuite>*.
- For Windows, this location is: C:\InCharge*<n>**<productsuite>*.

The <n> represents the EMC Smarts software platform version number. The <productsuite> represents the EMC Smarts product suite to which the product belongs. For example, on UNIX operating systems, EMC Smarts IP Availability Manager is, by default, installed to: /opt/InCharge7/IP/smarts. On Windows operating systems, this product is, by default, installed to: C:\InCharge7\IP\smarts. This location is referred to as BASEDIR/smarts.

Optionally, you can specify the root of BASEDIR to be something other than /opt/InCharge7 (on UNIX) or C:\InCharge7 (on Windows), but you cannot change the <productsuite> location under the root directory.

For more information about the directory structure of EMC Smarts software, refer to the *EMC Smarts System Administration Guide*.

EMC Smarts IP Management Suite products

The EMC Smarts IP Management Suite includes the following products:

- EMC Smarts IP Availability Manager
- EMC Smarts IP Discovery Manager
- EMC Smarts IP Performance Manager
- EMC Smarts IP Server Performance Manager
- EMC Smarts Adapter for HP OpenView NNM
- EMC Smarts Adapter for IBM/Tivoli NetView

Related documentation

In addition to this document, EMC Corporation provides a Help system for command line programs as well as product 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.

EMC Smarts documentation

Readers of this guide may find the following related documentation helpful. It can be found in the BASEDIR/smarts/doc/pdf directory.

Note: These documents are updated periodically. Electronic versions of the updated manuals are available on the Powerlink website: `http://Powerlink.EMC.com`.

- *EMC Smarts Documentation Catalog*

- *EMC Smarts System Administration Guide*
- *EMC Smarts ICIM Reference*
- *EMC Smarts Common Information Model Infrastructure Models Chart*
- *EMC Smarts Common Information Model Application/Business Models Chart*
- *EMC Smarts ASL Reference Guide*
- *EMC Smarts Perl Reference Guide*
- *EMC Smarts Dynamic Modeling Tutorial*
- *EMC Smarts MODEL Reference Guide*

EMC Smarts IP Management Suite documentation

The following EMC Smarts documents are relevant to users of the IP Management product suite:

- *EMC Smarts IP Management Suite Installation Guide*
- *EMC Smarts IP Management Suite Deployment Guide*
- *EMC Smarts IP Management Suite Configuration Guide*
- *EMC Smarts IP Management Suite Discovery Guide*
- *EMC Smarts IP Availability Manager User's Guide*
- *EMC Smarts IP Performance Manager User's Guide*
- *EMC Smarts IP Adapters User's Guide*
- *EMC Smarts IP Availability Manager Extension for IPv6 User's Guide*
- *EMC Smarts IP Availability Manager Extension for NAS User's Guide*
- *EMC Smarts Quick Start Guide*

Refer to the *EMC Smarts Documentation Catalog* for documentation resources provided with other EMC Smarts product suites.

Suggestions for searching PDF files

You may search across multiple PDF files using the Adobe Acrobat Reader software.

1. If the documentation is not accessible to all users of the EMC Smarts product suite, copy the contents of the *BASEDIR/smarts/doc/pdf d*irectory to a central location, such as a shared drive on your LAN, that operators and others may view.
2. To search throughout the documentation library, open the Acrobat Reader software.
   1. Choose **Edit** > **Search**, and enter a word or phrase.
   2. In the **Where would you like to search** option, select **All PDF Documents in** and type the pathname of the location where the PDF documents reside.

If you have more than one EMC Smarts product suite installed, you can set up cross-product document searches by copying files from the BASEDIR/smarts/doc/pdf directory for each product suite into this common documentation directory path.

Conventions used in this guide

EMC uses the following conventions for notes and cautions.

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

IMPORTANT

An important notice contains information essential to the operation of the software.

CAUTION

A caution contains information essential to avoid data loss or damage to the system or equipment. The caution may apply to hardware or software.

Typographical conventions

EMC uses the following conventions in this guide:

Normal font	In running text: • Interface elements (for example, button names, dialog box names) outside of procedures • Items that the user selects outside of procedures • Java classes and interface names • Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, filenames, functions, menu names, utilities • Pathnames, URLs, filenames, directory names, computer names, links, groups, service keys, file systems, environment variables (for example, command line and text), notifications
Bold	In procedures: • Names of dialog boxes, buttons, icons, menus, fields, and maps • Selections from the user interface, including menu items, and field entries • Key names • Window names • Code examples • Commands • User interface tabs
Italic	Used for: • Full publications titles referenced in text In procedures: • Directory paths • Filenames • Scripts • Important/unique words • URLs • Utilities • Variables
`Courier`	Used for: • System output • User entry
< >	Angle brackets used to enclose parameter or variable values supplied by the user.
[ ]	Square brackets used to indicate optional values.
|	Vertical bar used to indicate alternate selections. The bar means “or”.
{ }	Braces used to indicate content that you must specify (that is, x or y or z).
...	Ellipses used to indicate nonessential information omitted from the example.

%	Used to indicate a C shell prompt.
#	Used to indicate a C shell superuser prompt.
▼▲	Used to indicate a command is wrapped over one or more lines. The command must be typed as one line.

Pathname conventions

Directory pathnames are shown with forward slashes (/). Users of the Windows operating systems should substitute back slashes (\) for forward slashes.

Graphical conventions

If there are figures illustrating consoles in this document, they represent the consoles as they appear in Windows. Under UNIX, the consoles appear with slight differences. For example, in views that display items in a tree hierarchy such as the Topology Browser, a plus sign displays for Windows and an open circle displays for UNIX.

Smarts Manager

Unless otherwise specified, the term Smarts Manager is used to refer to EMC Smarts programs such as Domain Managers, Global Managers, and adapters.

Where to get help

EMC support, product, and licensing information can be obtained as follows:

Product information — For documentation, release notes, software updates, or for information about EMC products, licensing, and service, go to the EMC Powerlink™ website (registration required) at:

http://Powerlink.EMC.com

Technical support — For technical support, go to EMC Customer Service on Powerlink. To open a service request through Powerlink, you must have a valid support agreement. Please contact your EMC sales representative for details about obtaining a valid support agreement or to answer any questions about your account.

Sales and customer service contacts — For the list of EMC sales locations, please access the EMC home page at:

http://EMC.com/contact

1

Overview of the EMC Smarts Deployment Process

This chapter provides an overview of the deployment process for the portions of the EMC Smarts architecture that monitor and diagnose network infrastructure availability and performance.

- EMC Smarts architecture 1-2
- The deployment process 1-4
- Before-you-begin checklist 1-7

EMC Smarts architecture

This guide describes the deployment process for the following components in the EMC Smarts IP Management Suite and the EMC Smarts Service Assurance Management Suite:

- The EMC Smarts Discovery Manager (IP Discovery Manager) monitors and collects data for devices in a network. The IP Discovery Manager integrates the data but does not perform root-cause analysis.
- The EMC Smarts IP Availability Manager (IP Availability Manager) and EMC Smarts IP Performance Manager (IP Performance Manager) discover and monitor data for their corresponding domains. These managers then integrate and correlate the data to perform root-cause analysis for the devices in the related domain.
- The Service Assurance Manager correlates topology and events from various sources, including the Domain Managers and certain Adapters. The Service Assurance Manager also builds a graphical representation of these relationships in a topology map for display on the Global Console.
- The Global Console provides the interface for configuration and administration of EMC Smarts components and displays the root-cause of any failures using notifications and the topology map.
- The Service Assurance Manager Adapters (SAM Adapters) and the Adapter Platform provide additional sources of data when monitoring a network. The Adapter Platform normalizes and consolidates data from certain adapters before it is passed to Service Assurance. This guide covers deployment of the Adapter Platform and two of the adapters that send data to the Adapter Platform: the SNMP Trap Adapter (Receiver) and the Syslog Adapter.

Figure 1-1 "EMC Smarts components described in this guide" on page 1-3 shows the EMC Smarts components covered in this guide.



Figure 1-1 EMC Smarts components described in this guide

The deployment process

Figure 1-2, "The EMC Smarts deployment process," shows the EMC Smarts deployment process divided into four distinct phases:

- Phase 1: Designing the EMC Smarts deployment
- Phase 2: Installing and configuring the EMC Smarts components
- Phase 3: Validating the EMC Smarts deployment
- Phase 4: Tuning and maintaining the EMC Smarts deployment to improve performance

During each phase of deployment, it is vitally important that you document all aspects of the deployment that could be required to recreate, troubleshoot, and reconfigure the installation.



Figure 1-2 The EMC Smarts deployment process

Phase 1: Designing the EMC Smarts deployment

Designing an EMC Smarts deployment consists of gathering necessary information and then using the information to develop a plan for the EMC Smarts deployment.

Gathering the information is a process that involves both the designer and the network administrators. Network administrators must provide details of their network's OSI model Layer 2 and Layer 3 infrastructure:

- List the quantities and types of devices such as routers, switches, hubs, and bridges and their ports and interfaces. Include plans for adding or removing equipment during the time period that the deployment design will cover.
- Provide IP addresses for use in discovery, including seed systems and filtering.
- Describe network geography, including locations of Network Operations Centers (NOCs) and equipment.
- Specify LAN and link speeds throughout the network and their relationship to the network geography.

This information is then applied, using sizing guidelines, to develop a deployment design that will properly support the network. The design should also account for projected growth of the network.

The design is detailed in a Solution Architecture Diagram and a Deployment Build Guide. In their preliminary stages, the diagram and guide could be used in the answer to an organization's Request for Proposal (RFP) or Request for Quote (RFQ).

Once a deployment is contracted, the information in the Solution Architecture Diagram can be refined and verified through meetings with network administrators. If possible, using the IP Discovery Manager to inventory the network can provide a highly accurate starting point for sizing the deployment.

Design is covered in the following chapters:

- "Gathering Information for Designing" on page 2-1 (normally performed by systems engineers)
- "Designing the EMC Smarts Deployment" on page 3-1 (normally performed by systems engineers and either EMC Professional Services or a partner)
- "Planning for Discovery" on page 4-1 (normally performed by EMC Professional Services or a partner)
- "Designing Polling and Thresholds" on page 5-1 (normally performed by EMC Professional Services or a partner)
- "Deploying Syslog Processing" on page 6-1 (optionally performed by EMC Professional Services or a partner)
- "Designing for Administration of EMC Smarts" on page 7-1 (optionally performed by EMC Professional Services or a partner)

Phase 2: Installing and configuring the EMC Smarts components

Once the design is complete and has been reviewed by EMC Professional Services or an appropriately certified EMC partner, the next phase is installing and configuring components.

With any deployment, installation and configuration of EMC Smarts components are usually performed in stages. Each installed and configured segment is validated individually as described in the next phase of the deployment process. This phased approach eases troubleshooting.

Many organizations have specific procedural requirements that must be met before and during installation of new software products in their production environments. These requirements might include lab installations with performance validations and preproduction deployments. Lab configurations usually require the use of a testbed that is configured and equipped similarly to the production environment. Acceptance tests may be performed before the deployment to the production environment. After the production deployment, the lab or testbed may be used to test upgrades and, if required, patches.

Though this guide cannot cover organization-specific requirements, it does provide guidelines that may aid you in responding to these conditions.

Installing and Configuring EMC Smarts is covered in "Deploying EMC Smarts" on page 8-1. The deployment information, including software locations and all configuration choices should be recorded in the Deployment Build Guide. Normally, this phase is performed by the purchaser of the EMC Smarts deployment and either EMC Professional Services or a partner.

Phase 3: Validating the deployment

Validating the deployment ensures that all installed components are operational and can communicate with each other as required and that the appropriate components can properly discover and poll the network.

Usually logical segments are installed and validated to ease troubleshooting. Once all individual segments are installed and validated, the complete deployment must be validated from end to end. Included in this overall validation could be acceptance tests that demonstrate the functionality of the installation. Criteria for acceptance tests should be defined during the design phase. The execution of these acceptance tests and the results are then included in an installation or build report.

Deploying EMC Smarts components is covered in "Validating Your Deployment (Acceptance Testing)" on page 9-1. Normally, this phase is performed by the purchaser of the EMC Smarts deployment and either EMC Professional Services or a partner.

Phase 4: Tuning and maintaining the deployment

Once EMC Smarts is deployed and validated, you must ensure that it is operating at an optimal level. Tuning is the process of adjusting the configuration to improve performance. Note that this initial tuning process does not include rules writing or related maintenance: EMC Smarts does not require this type of maintenance.

This process can only be performed after *all* components are installed and validated to avoid inaccurate tuning.

The process of tuning an EMC Smarts deployment is covered in "Tuning Your Deployment to Improve Performance" on page 10-1. Normally, this phase is initially performed by either EMC Professional Services or a partner and the knowledge is transferred to the purchasing organization's staff. As the network grows and changes, the organization's staff will take on the task of tuning the EMC Smarts deployment to deal with the network changes. If the network changes are extensive, the original design may no longer be sufficient and redesigning the deployment may be required.

Before-you-begin checklist

Before you begin an EMC Smarts deployment, you must meet the requirements described in the Table 1-1. Each chapter in this guide includes a checklist. For ease of use, the checklists are all grouped together in "Design and Deployment Checklists" on page C-1.

Table 1-1 Before-you-begin checklist

Complete	Requirement	Description
❐	Possess an understanding of the EMC Smarts architecture and capabilities.	At a minimum, you must understand the concepts and EMC Smarts architecture described in the following documents: • *EMC Smarts IP Availability Manager User's Guide* • *EMC Smarts IP Management Suite Discovery Guide* • *EMC Smarts IP Management Suite Configuration Guide* • *EMC Smarts IP Performance Manager User's Guide* • *EMC Smarts Service Assurance Manager Configuration Guide* • *EMC Smarts Service Assurance Manager Introduction* • *EMC Smarts Service Assurance Manager Adapter Platform User's Guide* • *EMC Smarts System Administration Guide* • *EMC Smarts IP Management Suite Installation Guide* that accompanied your software product suite
		To improve your understanding, attend EMC Smarts training courses. Typically, deployment requires the knowledge equivalent to what is provided in the training courses on: • EMC Smarts IP Management Suite • EMC Smarts Service Assurance Manager • EMC Smarts Service Assurance Manager Adapter Platform
❐	Obtain contact information for the deployment team.	The contact list should include titles, responsibilities, and contact methods for all team members.
❐	Get nondisclosure requirements and negotiate an agreement.	Be aware of the requirements of the non-disclosure agreements that are in place for the deployment.
❐	Develop schedules and set milestones for early deliverable.	Scheduling a software deployment varies based on the size and scope of the deployment and the organization's requirements. Typical milestones might include: • Initial project meeting to define the deployment scope • Purchase of EMC Smarts software • Project development begins • Installation in test environment complete • Testing complete • Installation in production environment complete • EMC Smarts goes live Additional information on scheduling is beyond the scope of this guide.

2

Gathering Information for Designing

This chapter describes the first step in designing a deployment: Gathering the organizational and network-related information that is required to develop a successful EMC Smarts design.

- Determine the organization's requirements 2-2
- Obtain network information 2-3
- Determine number of managed network devices 2-6
- Gather network security information 2-11
- What other network features affect EMC Smarts? 2-12
- Architectural information checklist 2-13

Determine the organization's requirements

The design of an EMC Smarts deployment must support an organization's needs. Most Requests for Information (RFIs), Requests for Proposal (RFPs), or Requests for Quote (RFQs) begin with a description of the overall organization and its vertical market. Understanding the organization and its market can aid in making design choices. Typical vertical markets are listed in Table 2-1.

Table 2-1 Typical vertical markets

Typical Vertical Markets			
Communications E-Business Education	Financial General Business Government/Defense	Health Care Hosting Service Provider Network Outsourcers	Retail Transportation Wireless

Use an understanding of the business expectations in a vertical market to ensure a successful design. For example, each industry varies in the amount of downtime it can tolerate. As a general rule, financial organizations tolerate less downtime than organizations in the education vertical market. This can guide your design of polling and of escalation.

Obtain network information

Gather information that indicates the size of the network and how it is utilized to accomplish an organization's goals. A primary source for some of this information is an organization's RFI, RFP, or RFQ for the EMC Smarts deployment. The RFQ will normally contain details about a network's size and structure and deployment needs. Other sources of information include network diagrams and discussions with network administrators.

Obtain network diagrams

"Typical international network diagram" on page 2-4 shows a typical network diagram. To aid in design, the network diagram should include the physical geography of the network including locations for the following:

- Network Operations Center (NOC) and LANs
- All routing, bridging, and switching devices
- Firewalls
- Lower speed WAN links such as T1 links
- Higher speed network technologies such as FDDI and Gigabit Ethernet

Important IP addresses and address ranges should be listed on the diagram. Device names and device naming conventions should also be included.

Naming conventions

Device, interface, and port naming conventions are vitally important when adding customized processing to the deployment. Writing scripts to perform processing such as discovery post processing or advanced trap integration may only be practical in networks with naming conventions.



Figure 2-1 **Typical international network diagram**

Network priorities

Organizations reflect their business priorities in their network organization. Certain parts of the infrastructure are more important and must be monitored more closely to ensure availability. This can affect polling cycles and threshold groups. Once again, the organization's RFP should describe their needs and limitations. Make notes on the network diagram to indicate the priorities that you uncover.

Identify the types of equipment in the network

Create a list of all types of managed devices in your network and then compare the list with the EMC Smarts Certified Device list. This list is included with the software.

Any SNMP-capable device will be discovered and monitored by EMC Smarts using generic SNMP MIB II instrumentation if it is not yet certified by EMC Corporation: more specific instrumentation will not be available until the device is certified. Many uncertified devices can be field certified or, to provide the highest level of compatibility, certified in EMC Corporation labs.

To certify a device, EMC Corporation requires the following information:

- A MIB walk of the uncertified device using the IP Management Suite utility, **sm_snmpwalk**.
- Device containment information such as the IP addresses, number of cards, ports, and interfaces.
- A diagram showing how the device is connected to other devices in your network.
- MIB walks of all connected devices using **sm_snmpwalk**.

Other information specifically related to the use of the uncertified device in your network may also be required.

EMC Corporation periodically releases SmartPacks for the EMC Smarts IP Management Suite that provide certifications of new devices. The device certifications are intended to be as extensive as possible, but they may be prioritized to meet the needs of the majority of EMC Smarts deployments.

Determine requirements for installing software

Most organizations define some criteria for installing new software on their network. At a minimum, this might include software testing requirements.

In addition to responding to these requirements, plan on using a staging area before deploying. By staging the deployment, you can maintain a “clean” software distribution that does not include unwanted files, changes, and logs.

Integrating existing software with EMC Smarts software

Integrating third-party software products with EMC Smarts software requires the use of EMC Smarts adapters. Most integrations do one of the following:

- Exchange information. Determine the exact type of information (topology and/or events) that must be exchanged with any third-party software products.
- Allow access to the third-party software from the Global Console. Some adapters provide this capability as the Server Tools functionality.

EMC Corporation has developed many adapters that may already support the integration requirements of your deployment. Itemize the integration requirements so that the suitability of existing EMC Smarts adapters can be assessed by EMC Professional Services. This guide covers only the integration capabilities provided by the SNMP Trap Adapter (Receiver) and Syslog Adapter.

There are also extensive third-party software integration capabilities supported by other EMC Smarts adapters.

Determine number of managed network devices

Determine the number of network devices that will be managed by the EMC Smarts software deployment for two distinct purposes:

- The quantity of both ports and interfaces is vitally important for accurately sizing the EMC Smarts deployment. Sizing determines the number of Domain Managers to install and the required hardware configuration for the platforms where these managers will be installed.
- The quantity of systems is needed for EMC Smarts volume licensing. The discovery process will halt if the number of discovered systems or network adapters exceeds the licensed quantity.

In most cases, an organization's RFP or PFQ will estimate these quantities. This is the most common method of determining network devices: use the numbers from the RFP or RFQ as a starting point, and then improve the accuracy of those numbers by using common sense and by consulting with network administrators and other knowledgable personnel.

The most accurate method for obtaining the quantity of network devices is to use a Discovery Manager to inventory the network. Discovery Manager can be purchased with a travelling license so that it can be installed on a laptop PC and then used to autodiscover an organization's network via a temporary connection. For more information on using Discovery Manager, see "Using Discovery Manager to Obtain Network Information" on page A-1.

Refining the estimate for sizing an EMC Smarts deployment

The network size determines the time it takes to complete discovery, the memory required, and the server hardware that should be selected to support Domain Managers. Estimates of network size from network administrators are usually based on one of the following quantities:

- Quantity of ports and interfaces
- Quantity of routers and switches
- Quantity of total devices

Ultimately, the most important quantity for sizing the server hardware in a deployment is the total number of managed ports and interfaces. Managed ports and interfaces are monitored using SNMP and ICMP to determine status and connectivity.

After numerous deployments, EMC Corporation has developed ratios that help estimate the number of managed ports and interfaces with a high degree of accuracy. These ratios provide conservative estimates, which is especially important in large and very large networks. Being conservative is imperative, because EMC Corporation has determined that even in well-managed large networks, administrators can underestimate the number of ports and interfaces by 30% or more. Always use the multipliers in conjunction with validation, common-sense, and discussions with network administrators to improve the accuracy of the estimates.

Estimates based on ports and interfaces

The most accurate method of determining network size is by retrieving the count of ports and interfaces from a network discovered by the IP Discovery Manager.

If the network includes virtual routers, the count will include their interfaces in addition to the physical ports and interfaces. A virtual router is a software emulation of a router implemented within a physical router or switch.

Calculate the estimated number of managed ports and interfaces using the ratios of 90% managed for interfaces and 5% managed for ports:

Estimated number of managed ports and interfaces	=	90% Interfaces	+	5% Ports

For example, if the network has 2,300 interfaces and 18,000 ports, then calculate the managed ports and interfaces as follows:

Estimated number of managed ports and interfaces	=	90% Interfaces	+	5% Ports
	=	((90/100) * 2,300)	+	((5/100) * 18,000)
	=	2,070	+	900
	=	2,970		

All values that were less than one were rounded up in these calculations.

Estimates based on routers and switches

When the number of ports and interfaces is not available, the next best method is to use the number of routers and switches to make an estimate:

- Obtain the total number of routers and the total number of switches in the network. Include virtual routers in this count of physical systems if they are used in the network. EMC Corporation has developed two multipliers to represent the typical number of ports and interfaces for these devices: 25 interfaces per router and 60 ports per switch. These values are typical in most deployments.

  Total ports = 60 * switches

  Total interfaces = 25 * routers

- Using the total ports and total interfaces, estimate the number of managed ports and interfaces. Use the managed ratios for ports (5%) and interfaces (90%) to calculate managed ports and interfaces

  Managed ports = 5% total ports

  Managed interfaces = 90% total interfaces

- Add 30% to the estimate of managed ports and interfaces for uncertainty.

  Estimated managed ports and interfaces = (managed ports + managed interfaces) + 30%

- For example, if a network has 95 routers and 305 switches, the estimate for the number of managed ports and interfaces is calculated as follows:

Total ports = 60 * switches
= 60 * 305
= 18,300

Managed ports = 5% total ports
= (5/100) * 18,300
= 915

Total interfaces = 25 * routers
= 25 * 95
= 2,375

Managed interfaces = 90% total interfaces
= (90/100) * 2,375
= 2,138

Estimated managed ports and interfaces = (managed ports + managed interfaces)+ 30%
= (915 + 2,138) + 30%
= 3,053 + ((30/100) * 3,053)
= 3,053 + 916
= 3,969

All values that were less than one were rounded up in these calculations.

Estimates based on devices (Level 2/Level 3)

The least accurate method for estimating the size of a deployment uses only the total of level 2 and level 3 devices:

- Use the total of devices, including virtual routers, to estimate the number of ports and interfaces in the network. EMC Corporation has developed the multiplier of 50 to represent the typical number of ports and interfaces per device. This multiplier assumes a split of about 30% routers and 70% switches in the network.

Total ports and interfaces = 50 * devices

- Using total number of ports and interfaces, estimate the number of managed ports and interfaces. Assume that 35% of the total ports and interfaces are managed.

Managed ports and interfaces = 35% total ports and interfaces

- For example, if a network has 400 level 2 and level 3 devices, the estimate for the managed ports and interfaces is calculated as follows:

Total ports and interfaces	= 50 * devices
	= 50 * 400
	= 20,000

Managed ports and interfaces	= 35% total ports and interfaces
	= (35/100) * 20,000
	= 7,000

All values that were less than one were rounded up in these calculations.

Note that due to the lack of specific information, this number is probably less accurate than the numbers produced by other methods and should always be discussed with network administrators.

Accounting for sub-interface monitoring

EMC Smarts permits performance monitoring of sub-interfaces. If you intend to monitor the sub-interfaces using the IP Performance Manager, each sub-interface must be counted as an interface and added to your total of managed ports and interfaces.

The quantity of sub-interfaces is very difficult to estimate: for example, some Cisco devices permit you to configure thousands of sub-interfaces. If your network uses technologies such as frame relay, ATM, or ISDN, pay particular attention to the sub-interface configurations. Consult network administrators to determine a reasonable maximum quantity of sub-interfaces.

Accounting for network growth

Your deployment design should account for network growth over the expected life of the design. The network growth rate will relate to the vertical market environment and the organization's plans, so the organization must provide you with growth estimates.

Total managed ports and interfaces (P&I)	=	**Current managed P&I**	+	**Additional managed P&I due to growth**

Network growth planning can involve very complex calculations, but this is beyond the scope of this guide. For illustrative purposes, a simple method based on percentages is used here. For example, if a network with 5500 managed ports and interfaces was expected to grow by 10% per year, to calculate the size after one year:

Total managed ports and interfaces (P&I)	=	**Current managed P&I**	+	**Additional managed P&I due to growth**
	=	5,500	+	(10/100) * 5,500
	=	5,500	+	550
	=	6,050 after one year		

To continue this example, if the design life is two years, then recalculate to add an additional 10% growth of the deployment over the second year:

Total managed ports and interfaces (P&I)	= 6,050	+ (10/100) * 6,050
	= 6,050	+ 605
	= 6,655 after two years	

For the same network and two year design life, with an expected growth rate is 20%, the calculations are as follows:

Total Managed ports and interfaces (P&I)	**= Current managed P&I**	**+ Additional managed P&I due to growth**
	= 5,500	+ 20% (5,500)
	= 5,500	+ 1,100
	= 6,600 after one year	
	= 6,600	+ 20% (6,600)
	= 6,600	+ 1,320
	= 7,920 after two years	

All values that were less than one were rounded up in these calculations.

Determine quantities of devices for licensing

EMC Smarts uses volume licensing to determine the size of the topology that Availability Managers are permitted to discover and manage in the deployment.

Volume licensing counts either or both of the following:

- All managed network adapters (ports and interfaces) in the topology.
- All managed systems in the topology. Systems include routers, switches, hubs, virtual routers, security devices, hosts, servers, desktop or laptop PCs, workstations, probes, terminal servers, printers, IP phones, wireless access points (WAPs), and CSUs/DSUs.

An accurate count of systems ensures that the deployment can manage all the devices appropriately. The discovery process will halt if the number of discovered systems or network adapters exceeds the licensed quantity.

As the network grows and more systems are added, you can obtain additional volume licenses from EMC Corporation.

Gather network security information

Determine the level of security for the network that the EMC Smarts software will monitor so that it can be configured to a corresponding level of security. For example, the security needs of a network in a financial, defense, or health care vertical market may be greater than in the manufacturing vertical market. Enumerate security preferences, such as the use of passwords, encrypted password storage, and encrypted communications to guide you when configuring EMC Smarts security capabilities.

There are many security-related network features that will affect the deployment. These include:

- Firewalls between parts of the deployment. Appropriate EMC Smarts components must be able to poll the network, receive traps, and communicate with other EMC Smarts components. Certain TCP and UDP ports will need to be opened in the firewalls to facilitate these communications.
- Use of access lists. If access lists are used, the IP addresses of servers running EMC Smarts products must be added to the access list of devices that will communicate with EMC Smarts. EMC Smarts must have full access to browse the MIBs of the devices.
- Use of SNMP versions and their respective security capabilities. The version of SNMP used to communicate with the network devices can provide dramatically different levels of security. With SNMP V1 or V2C, the security is provided through the use of SNMP community strings. To properly configure EMC Smarts, you must know the SNMP read community strings for all SNMPV1 and SNMP V2C devices that will be managed.

  For communications to devices using SNMP V3, the requirements are much greater. Obtain the values for these configuration parameters for each SNMP V3 device:

  - SNMP V3 user name
  - SNMP engine ID (Optional. If wrong or omitted, discovery will find it.)
  - Authentication protocol (Only MD5 and SHA are supported. MD5 is the default.)
  - Authentication password (if Authentication protocol is used)
  - Privacy protocol (AES and DES are supported. NONE is the default.)
  - Privacy password (if Privacy protocol is used)
  - Context name (if used)

What other network features affect EMC Smarts?

Other network features may affect the deployment design. Consider the following questions when gathering information on the network:

- Does the organization require failover capabilities in network software?
- Is there an "out-of-band" network just for management information? For example, are certain ethernet ports just for management information? This is typical with some deployments in the financial and military/defense vertical markets.
- Who will use the Global Console and what are their needs? Who will be the operators and administrators? What are their access privileges? What network availability and performance information do they need to view?
- What are the issues related to network latency, bandwidth, and speed available for network management traffic?

Architectural information checklist

Use the following checklist to aid in gathering information for your architectural design. The checklist includes space for writing information; record the information here or in your design documentation. Each chapter in this guide includes a checklist. For ease of use, the checklists are all grouped together in "Design and Deployment Checklists" on page C-1

Table 2-2 Architectural Information Checklist

Complete	Task	Descripiton	Related documentation
❑	Describe the organization's requirements and expectations.	Organization's vertical market: _____ (Reference to an organization's documentation) _____ _____ _____	"Determine the organization's requirements" on page 2-2
❑	Obtain network diagrams.	Ensure the diagrams include the locations of the following: • Network Operations Center (NOC) and LANs • Routing and switching devices • Firewalls • WAN links • High speed network technologies such as FDDI and Fast or Gigabit Ethernet In addition, important IP addresses and address ranges should be indicated.	"Obtain network diagrams" on page 2-3
❑	If possible, schedule and discover the network.	If you intend to use the IP Discovery Manager to inventory the organization's network, schedule a time to perform the process and then discover the network as scheduled.	"Obtain network information" on page 2-3
❑	Describe the organization's network priorities.	Document these priorities in the Deployment Build Guide.	"Network priorities" on page 2-4
❑	Get the organization's testing/acceptance requirements.	Your design may be required to meet test and acceptance requirements. Obtain any specifications that cover integration testing, user acceptance testing, and operational acceptance testing. You may be required to write an installation or deployment report that follows an organization's particular standards.	"Determine requirements for installing software" on page 2-5
❑	Describe the organization's requirements for installing new software.	❑ Lab installation and testing ❑ Staging (*strongly* recommended) ❑ Preproduction deployment ❑ Shadow operation period (existing MoM still used) ❑ Other _____ Document these requirements and how the design meets them in the Deployment Build Guide.	"Determine requirements for installing software" on page 2-5
❑	List the products that currently monitor the network and will be integrated with EMC Smarts.	EMC Smarts' open architecture allows easy integration with third-party software. Many networks have at least a rudimentary network availability monitoring or a legacy system such as HP OpenView or IBM NetView. Document the products (including version) in deployment Build Guide.	"Integrating existing software with EMC Smarts software" on page 2-5
❑	List device types to manage.	To ensure that devices are certified, obtain a list of the manufacturers and models for all devices in the network. Document the types of managed devices in the Deployment Build Guide.	"Identify the types of equipment in the network" on page 2-5
❑	Determine the number of managed ports and interfaces in the network.	Document all quantities and calculations used to determine the number of managed ports and interfaces in the Deployment Build Guide.	"Determine number of managed network devices" on page 2-6

Table 2-2 Architectural Information Checklist (continued)

Complete	Task	Descripiton	Related documentation
❑	Estimate potential growth in quantity of managed devices.	The deployment must support potential network growth. Estimate the growth over a specific time period. Document the calculations in the Deployment Build Guide.	"Accounting for network growth" on page 2-9
❑	Estimate number of managed systems and network adapters for licensing.	The deployment can only discover and manage the quantity of systems and network adapters that are licensed. Document the quantities in the Deployment Build Guide.	"Determine quantities of devices for licensing" on page 2-10
❑	Describe the network security.	Describe security features such as the firewalls that will be between parts of the deployment and if access lists are used. Obtain SNMP security parameter values for each device where they are used: for SNMP V1 and V2C, obtain community strings; for SNMP V3, obtain the user name, SNMP engine ID (optional), authentication protocol and password, privacy protocol and password (currently EMC Corporation supports AES and DES privacy protocols), and context name, if used. Document the security features in the Deployment Build Guide.	"Gather network security information" on page 2-11
❑	List any other network requirements or features that may affect EMC Smarts.	Document the features in the Deployment Build Guide.	"What other network features affect EMC Smarts?" on page 2-12

3

Designing the EMC Smarts Deployment

This chapter provides guidelines for producing an initial design of the EMC Smarts deployment using the organization and network information that you have gathered. This design usually takes the form of a Solution Architecture Diagram and a Deployment Build Guide.

- Documenting the deployment 3-2
- Determine the required size of the EMC Smarts deployment 3-3
- Partitioning networks 3-5
- Adding information to the Solution Architecture Diagram and Deployment Build Guide 3-6
- Locating EMC Smarts managers and platforms 3-7
- Security and licensing considerations 3-9
- Designing for overlapping (duplicate) IP networks 3-11
- Designing acceptance tests 3-12
- Solution Architecture Diagram checklist 3-13

Documenting the deployment

The most useful way to document the design of your EMC Smarts deployment is to create a Solution Architecture Diagram and record implementation details in a Deployment Build Guide.

The Solution Architecture Diagram

Based on the complexity of the deployment, a Solution Architecture Diagram may actually be a set of diagrams documenting various levels of the architecture.

The diagram relates both physical and logical choices for your EMC Smarts architecture in an easily understood manner. This diagram graphically reflects your design choices and will be an important part of the review and approval process for your design.

The Solution Architecture Diagram should always include:

- A logical representation of the EMC Smarts components that will be installed
- Locations for each EMC Smarts component including the name and IP address of the host and the geographical location of the host
- Connections between EMC Smarts components and the ports that are used for communications
- Connections, including port numbers, between EMC Smarts and external sources such as networks and third-party software products

This chapter describes how to start the Solution Architecture Diagram. This diagram cannot be completed until the design is complete, so the design portion of the guide includes directions for adding information to the diagram.

The Deployment Build Guide

To record the specifics of the EMC Smarts deployment design and implementation, create a document called a Deployment Build Guide. As with the Solution Architecture Diagram, this chapter describes information that you should add to the Deployment Build Guide. The Deployment Build Guide should include the complete design and all installation specifications, validation results, and tuning activities.

Start the Deployment Build Guide by recording all the information that you have gathered on the network. Include a copy of the network diagram that you have already obtained. As you continue the deployment process, this guide will include recommendations for adding other information to the Deployment Build Guide.

Determine the required size of the EMC Smarts deployment

The size of a deployment is directly related to quantity of managed ports and interfaces that the deployment must support. Based on the values for managed ports and interfaces that you developed using the formulas in "Determine number of managed network devices" on page 2-6, you can begin to determine how many Availability Managers and Performance Managers are needed.

As a first step, work with the network administrators to choose the platform equipment tier that is appropriate for the organization's resources and personnel. For example, familiarity with the operating system or with a vendor's equipment may direct the choice. Typical hardware for the equipment tiers and operating systems is listed in Table 3-1.

Table 3-1 Typical hardware for the equipment tiers

Operating System	Platform Equipment Tier			
	Small (2-4 low end CPUs, 1 GB RAM) *	**Medium (4+ CPUs, 2 GB RAM)**	**Large (16+ CPUs, 4 GB RAM)**	**Extra Large (32+ CPUs, 8 GB RAM)**
Solaris	SunFire V210 or V240 (preferred)	SunFire V280	SunFire V280 with fastest CPU	SunFire V490 with fastest CPUs
HP-UX ("-x" indicates number of processors)	HP rp3410-2, rp3440-4, rp4410-4 or rp4480-8	HP rp7420-16 or rp8420-32	Superdome (from 16 to 64 processors)	Superdome (from 16 to 64 processors)
Linux	Any vendor, 1-2 Xeon 2.0 GHz	Any vendor, 2 Xeon 3.0 GHz	Any vendor, 2 Xeon 3.6 GHz	Any vendor, 4 Xeon 3.6 GHz
Windows				
* The small tier configurations are less than the minimum requirements listed in the EMC Smarts installation guides. These configurations are not recommended unless you are very experienced with EMC Smarts deployments and have very reliable network size information. Contact EMC Technical Support for the latest hardware equipment specifications.				

Note: SNMP V3 authorization protocol requires more computer resources than earlier versions of SNMP. If you are using a large number of SNMP V3 devices or systems, you will need to increase the tier size you select.

Advantages of larger platform equipment tiers

The larger the platform equipment tier, the larger the portion of the network that a single Domain Manager installed on that hardware can support. Advantages and disadvantages of using a single Domain Manager on larger platforms are listed in Table 3-2.

Table 3-2 Advantages and disadvantages of a single process on larger platforms

Advantage:	Disadvantage:
• Easier to maintain. • Split topology is not required. Splitting a topology into multiple domains is a complex process that must be updated as the network topology changes. • Saves bandwidth for discovery: separate Availability or Performance Managers must each discover the network. Multiple Availability Managers in a split topology must discover parts of the network twice. • If Availability and Performance Manager are installed as a single process, synchronizing the topology between them is not necessary.	• Scalability is limited. • Requires tuning. • Increased risk of encountering single-threaded bottlenecks. • The memory required by the process may reach or exceed operating system limits. • Requires expensive hardware.

After choosing the equipment tier that best suits the organizational requirements, determine how many Domain Manager are required. Dividing your estimate of the total managed ports and interfaces by one of the values in Table 3-3. Select the EMC Smarts functionality required: IP Availability Manager, Availability Manager/Performance Manager (single process), or IP Performance Manager.

When Availability Manager and Performance Manager are installed as a single process rather than two separate processes on the same platform, the single process supports a more limited number of network device ports and interfaces. This is caused by the additional attributes that Performance Manager must track per port or interface (10 times more than Availability Manager). Performance Manager also tracks five more attributes per system.

Table 3-3 Port and interface support by platform equipment tiers

Sizing Considerations		Platform Equipment Tier			
		Small	Medium	Large	Extra Large
Managed Ports and Interfaces (P&I) supported	AM only:	10,000	25,000	50,000	100,000
	AM/PM or PM only:	5,000	15,000	25,000	50,000
AM = Availability Manager; PM = Performance Manager					

For example, if EMC Smarts must support 30,000 managed ports and interfaces in an availability-only deployment, you would require:

- Three Availability Managers with each one on a small-tier platform
- Two Availability Managers with each one on a medium-tier platform
- One Availability Manager on a single, large-tier platform

Partitioning networks

To avoid processing bottlenecks such as reconfiguration in a very large network, it may be necessary to partition the network. Splitting a very large topology into multiple EMC Smarts domains is a complex process. EMC Professional Services can perform this process for any deployment.

Multiple EMC Smarts managers on a single platform

When splitting the topology is necessary or desired, you can install multiple Domain Managers on a single platform. Simply add the CPUs and memory requirements for each Domain Manager you wish to install on a platform to determine if the platform can support the configuration.

Adding information to the Solution Architecture Diagram and Deployment Build Guide

Include the appropriate number of IP Availability Managers and IP Performance Managers on the Solution Architecture Diagram. Enclose the managers in a box to represent the appropriate host as shown in "Typical start for Solution Architecture Diagram showing hosts and underlying domain managers." on page 3-6.



Figure 3-1 Typical start for Solution Architecture Diagram showing hosts and underlying domain managers.

For the Deployment Build Guide, document equipment choices and start a table to document components on each host.

Locating EMC Smarts managers and platforms

In choosing locations for the platforms that support Domain Managers, there may be restrictions on locations that are not related to network and product efficiency:

- Locating Domain Managers might be based on geographical requirements. Some organizations might require all Domain Managers to be based in a single Network Operations Center (NOC). Others might have the Service Assurance Manager in the NOC and the underlying Domain Manager located in regional data centers.
- Locating Domain Managers might be based on corporate organizational requirements. For example, organizations with distributed management might require that the deployment be partitioned to support a portion of a network that is split along bureaucratic rather than technical lines.
- Locating Domain Managers might be based on the network's security design. For example, if parts of the network are highly secured and SNMP or ICMP polling between these network segments is not allowed, separate IP Availability Managers would have to support each segment.

Polling and discovery are influenced by network speed and latency. If possible, consider network efficiency when locating Domain Managers. Avoid configurations that require EMC Smarts to discover or poll large portions of the network across lower speed WAN links or other network bottlenecks. Consider placing Domain Managers on higher speed LAN networks.

Add system names and Domain Manager names to your Solution Architecture Diagram. Define IP addresses and dedicated port numbers when needed.

Establish a host naming convention and an Domain Manager naming convention before you settle on any names. For example, a service provider using separate domains to support each of its customers chooses to use the name of the customer as part of the name of each Domain Manager that supported the customer.

Specify the locations for the hosts that support the IP Availability Manager and IP Performance Managers on the Architecture Solution Diagram as shown in "Adding locations to the Solution Architecture Diagram" on page 3-8.



Figure 3-2 Adding locations to the Solution Architecture Diagram

Security and licensing considerations

Considering volume licensing configurations

In certain cases, you may use multiple license servers in your network for the same reasons that you split a topology or choose different locations for Domain Managers: geographical requirements, corporate organizational requirements, or the network's security design. Using multiple license servers requires the corresponding number of EMC Smarts licenses.

When multiple license servers are used, each Domain Manager will use blocks of licenses from a specific license server. Therefore, the license blocks must be divided among multiple licenses and the appropriate license must be deployed to each license server. Then, the IP Availability Manager that relies on a specific license server will have the appropriate licenses to support the systems that it must manage.

Consult the Volume Licensing section of the *EMC Smarts System Administration Guide* for more information.

Considering security and firewalls

Based on the security information you obtained earlier, you must plan designing solutions so that EMC Smarts can function properly in the network's security environment:

- For communication between Domain Managers across firewalls, plan on opening a hole in the firewall for EMC Smarts communications. Certain UDP and TCP ports must be opened for proper communications:
  - SNMP Polls: Port 161
  - SNMP Trap: Port 162
  - Broker: Port 426
  - License Manager: Port 1744
  - Domain Manager: one port each, which can be configured.
  - EMC Smartsadapters, including SNMP Trap Adapter (Receiver) and Syslog Adapter: See "Deploying Syslog Processing" on page 6-1 and "Designing Trap Processing" on page 6-1.
- If access lists are used, plan on deploying the IP addresses of hosts that include Domain Managers to the access list of devices that will be managed. EMC Smarts must have full access to browse the MIBs of the devices. (The specific MIBs are listed in the *EMC Smarts IP Availability Manager User's Guide* and/or the *EMC Smarts IP Performance Manager User's Guide*.) Depending on the network size and complexity, this may require scheduling to obtain support from the organization's network personnel.
- You must have a listing of SNMP versions and related security parameter values that are used by specific devices in the organization's network. Due to security concerns, it may not be appropriate to include them in the Deployment Build Guide.

In addition, consider the level of security to configure for EMC Smarts products. The EMC Smarts security mechanisms support various levels of user authentication and both authentication and encrypted communication between EMC Smarts products. Ensure that you understand the capabilities described in the *EMC Smarts System*

Administration Guide and then choose a level of security that is appropriate for the deployment.

Considering high availability configurations

Failover is a complex configuration that currently requires aid from the EMC Professional Services organization. Consult with Professional Services if your installation requires this capability.

Designing for overlapping (duplicate) IP networks

EMC Smarts allows service providers who offer managed IP network services to centrally manage the private IP networks of customers who use identically-numbered IP address spaces. Enterprise users facing similar issues due to acquisitions of new networks can also use this ability. For more information, see "Managing Overlapping IP Networks With EMC Smarts" on page B-1.

Designing acceptance tests

Acceptance tests may be required for different portions of EMC Smarts functionality. Be aware of the requirements and develop acceptance criteria with the aid of the network administrators and other organization personnel. Include all necessary acceptance tests in the Deployment Build Guide for the deployment. At a minimum, you must develop completion criteria that the organization's project managers approve. Document these completion criteria in the Deployment Build Guide and use them in validation.

Solution Architecture Diagram checklist

Use the Solution Architecture Diagram to document your initial overall design of the deployment. Each chapter in this guide includes a checklist. For ease of use, the checklists are all grouped together in "Design and Deployment Checklists" on page C-1

Table 3-4 Solution Architecture Diagram checklist

Complete	Task	Description	Related documentation
❑	List important device quantities on the Solution Architecture Diagram and in the Deployment Build Guide.	Start the Solution Architecture Diagram by listing the totals for Routers, Switches, Hubs, Bridges, Hosts, Ports and Interfaces. Include expected growth rate and estimates for managed ports and interfaces. Also document the quantities in the Deployment Build Guide.	"Documenting the deployment" on page 3-2
❑	Choose a tier size for platforms supporting Domain Managers.	❑Small (P&I:10K/AM or 5K/AM PM) ❑Medium (P&I:25K/AM or 15K/AM PM) ❑Large:(P&I:50K/AM or 25K/AM PM) ❑Very Large:(P&I:100K/AM or 50K /AM PM) Document the choice in the Deployment Build Guide.	"Determine the required size of the EMC Smarts deployment" on page 3-3
❑	Determine the quantity of Domain Managers required.	Include representations on the Solution Architecture Diagram.	"Determine the required size of the EMC Smarts deployment" on page 3-3
❑	Document equipment supporting EMC Smarts components.	Document the hardware and operating system supporting each of the EMC Smarts components. Start a chart for each host in the Deployment Build Guide and list the EMC Smarts components on the host.	"Determine the required size of the EMC Smarts deployment" on page 3-3
❑	Locate the hosts supporting the EMC Smarts components.	Document choices on the Solution Architecture Diagram and in the Deployment Build Guide.	"Locating EMC Smarts managers and platforms" on page 3-7
❑	Determine license server and licensing configuration requirements.	Document requirements on the Solution Architecture Diagram and in the Deployment Build Guide.	"Considering volume licensing configurations" on page 3-9
❑	Determine security requirements.	Document requirements in the Deployment Build Guide.	"Considering security and firewalls" on page 3-9
❑	Is Failover Capability Required for EMC Smarts?	❑No ❑Yes: Contact EMC Professional Services. Document choices on the Solution Architecture Diagram and in the Deployment Build Guide.	"Considering high availability configurations" on page 3-10
❑	Determine if overlapping IP networks are used.	Document needs in the Deployment Build Guide.	"Designing for overlapping (duplicate) IP networks" on page 3-11
❑	Plan acceptance tests and completion criteria.	Document in the Deployment Build Guide as each portion of EMC Smartsfunctionality is designed. Use them in validation.	"Designing acceptance tests" on page 3-12

4

Planning for Discovery

This chapter provides guidelines for designing the discovery process for your deployment. Before beginning the design of your discovery process, read the comprehensive description in the *EMC Smarts IP Management Suite Discovery Guide*.

- Designing discovery 4-2
- Discovery design checklist 4-8

Designing discovery

Discovery is the process of creating a representation of the managed network topology within a Domain Manager's repository. Data is collected by ICMP and SNMP to create instances of the managed systems, their dependencies, and their internal components. Discovery is an expensive process in terms of required processing and network resources, so making the appropriate choices during design is imperative. The discovery implementation is very flexible and allows many variations.

To simplify discovery design, consider discovery in two separate stages:

- Initial topology discovery
- Maintenance and subsequent topology discovery

Initial discovery

The initial discovery of the network topology is usually the most time consuming and expensive in terms of network resources, so scheduling the initial discovery is very important. Accurately predicting the duration of the initial discovery is virtually impossible due to network variance, so assume the discovery will take a considerable period during this initial discovery. This is due to the time-consuming process of creating objects with defined relationships for every device in the topology.

Discover the network over a long period of low utilization, such as a weekend or holiday. Plan on this even if EMC Smarts is deployed in a test environment, because discovery can be adversely affected by heavy network utilization.

Other than scheduling, the most basic choice in the discovery design is whether or not to use autodiscovery. When autodiscovery is not used, a comprehensive seed file is required, or topology can be imported from a third-party source.

Using a comprehensive seed file

A seed file contains a list of IP addresses or system names to discover. This makes it the simplest to design, quickest to run, and provides the most control. However, to create a comprehensive seed file, the network must be well-documented and the network topology information must be complete and accessible.

For complete details on creating and using the seed file, refer to the *EMC Smarts IP Management Suite Discovery Guide*.

Importing topology using third-party software

With topology import, the IP Manager discovers a set of candidate systems imported from a third-party source through an EMC Smarts IP adapter. EMC Smarts can import topology information from many sources, including HP OpenView NNM and IBM/Tivoli Netview. This method requires an adapter such as the EMC Smarts Adapter for HP OpenView NNM or the EMC Smarts Adapter for IBM/Tivoli NetView described in the *EMC Smarts IP Adapters User's Guide*.

Topology import requires that at least one discovery filter be used. The filter or filters should be defined broadly enough to accept all possible IP addresses and system names provided by the third-party software.

For more on topology import see the *EMC Smarts IP Management Suite Discovery Guide*.

When should you use autodiscovery?

With autodiscovery, EMC Smarts automatically discovers your network from an agent or a seed file containing one or more seed systems. The discovered devices are probed for IP addresses of their neighbors, and the autodiscovery cycle continues until no more new IP addresses match the discovery filters.

During an initial discovery, using autodiscovery can be timesaving, particularly when topology information is incomplete as often happens when a network is constantly in flux.

Autodiscovery requires a partial seed file or the name of an agent device to begin the discovery. An appropriate agent device or system must be selected. Consider including a device or system enabled with Cisco Discovery Protocol (CDP), Extreme Discovery Protocol (EDP), or Foundry Discovery Protocol (FDP) to improve autodiscovery coverage. At the very least, use a device that is not at a network edge as an agent.

Autodiscovery also requires appropriate autodiscovery filters and configuration.

Though this method requires more resources than other methods, with appropriate discovery filtering, autodiscovery is very efficient and requires little in additional resources.

In addition, autodiscovery can be limited and controlled by specifying manual addition of discovered systems to the topology and setting an appropriate topology system limit.

When should you not use autodiscovery?

Using autodiscovery is not advised in the following deployment configurations:

- New network devices are constantly being phased in. During the phase in period, the devices are accessible on the network, although not fully operational and are being tested. Though you intend to eventually add the devices to the topology, adding the devices at this point would cause spurious notifications or fill the device pending list and obscure devices that should be discovered.
- Specific devices are accessible on the network but will never be managed. For example, ISPs might want to discover a router and its interfaces, but not the devices connected to the router from the ISP's client side. If the connecting devices do not follow a naming convention or can change without notice, it may be very difficult to define autodiscovery filters and exclude filters that ensure that the devices are not discovered or placed in the pending list.
- The network uses many SNMP read community strings. Currently, autodiscovery allows four and you can increase the number by reconfiguring *discovery.conf*, but in some situations dozens or more strings may be used and filtering becomes impractical. In addition, you may not want to use the community strings in a secure environment—devices may write community strings to syslog files when the devices are polled with strings that do not match.
- The network uses an inventory database and devices are being commissioned through the use of the database.
- The network uses SNMP V3 devices extensively. Autodiscovery cannot discover devices that use SNMP V3.

Using autodiscovery during initial discovery

Autodiscovery is particularly useful when topology information is incomplete. Even in network environments where autodiscovery is not recommended, autodiscovery can sometimes be used during an initial discovery to inventory the network and

create a comprehensive topology. After the initial discovery, autodiscovery could be disabled so that it is not triggered by subsequent full discoveries or pending list discoveries.

If you determine it is inappropriate to use autodiscovery on a regular basis, but want to take advantage of autodiscovery during the initial discovery, configure autodiscovery carefully. Ensure that the autodiscovery filters and the exclusion filters in the *discovery.conf* file are properly configured. Consider using the ***Ask before adding new systems*** option with each discovery filter to ensure the greatest control of the discovery process.

When the discovery process is complete, review the topology and discover appropriate devices on the device pending list. Remember that discovering devices on the pending list will trigger autodiscovery to restart. After completing discovery, create a comprehensive seed file using **sm_tpmgr --dump-agents**.

If necessary, disable autodiscovery after the initial topology discovery is complete.

Topology maintenance and subsequent discovery

After the initial discovery of the network topology, you must choose a schedule for subsequent discovery processes to maintain an accurate topology:

- Full discovery of the existing topology should be scheduled to occur at least once per week. Once again, scheduling full discovery during a long period of low network utilization such as a weekend is very important.

  For example, a large multinational bank discovers devices for a domain on Saturdays at 1:00 pm. Specific times are scheduled using **cron** or **sm_sched** to invoke **sm_tpmgr --discover-all** for the domain.

  The Global Console allows you to specify an initial date and time and a discovery interval, the time between the initiation of discoveries. If a discovery is in progress when it is time for the next discovery, the Topology Manager skips that full discovery and waits until the discovery interval elapses again to begin discovery. An exception is written to the IP Manager's log file.

  Note: This is different from the pre-6.3 behavior of the discovery interval which allowed you to specify a specific time.

- Pending discovery should be scheduled to occur at least once per day. This is the discovery of any devices that were placed in the Pending Devices list. Like full discovery, schedule the pending discovery during a period of low network utilization. The duration of a pending discovery is usually much shorter than full discovery, so schedule it during a relatively idle work shift.

  For example, a regional service provider discovers pending devices on weekdays at 2:00 am. As with full discovery, specific time intervals are scheduled using **cron** or **sm_sched** to invoke **sm_tpmgr --discover-pending** for the domain. The Global Console only allows you to specify an interval of time since the start of the last pending discovery was attempted. An exact time cannot be specified.

Note: If you use **cron** or **sm_sched** to schedule full and pending discoveries, uncheck the Enable Full Discovery option and use 99 days for Discover Pending Interval at the Domain Manager Administration Console.

Running the discovery processes more often will provide a more accurate network topology, but you must consider both your needs and the cost in terms of resources. If

your network changes more often than the recommended discovery schedules, then you must shorten the time between discoveries. Typical network changes include anything from hardware changes, such as adding cards or devices, to configuration changes such as reassigning IP addresses or modifying a VLAN.

Adding new systems to an existing topology

Remember that new devices will only be added to the topology if you do one of the following:

- Enable autodiscovery and create an appropriate filter configuration for automatic addition. New systems are NOT automatically found during discovery (or added) unless autodiscovery is enabled. If enabled, autodiscovery occurs *whenever* either a full discovery occurs or a pending device is successfully discovered. When full discovery takes place, all the devices already in the topology will be probed during discovery in an attempt to autodiscover new devices. If a topology is unstable, autodiscovery is a very useful feature.
- Use the add agent or import from seed file functions at the Global Console to discover new devices by name or IP address. Using comprehensive seed files is the typical way to add systems to networks with stable topologies that are well documented. Service providers who manage devices under contract will typically use this approach to avoid discovering and managing devices that they are not paid to manage.

These two methods of adding new devices can be used separately or combined. When combined, adding a system manually using a seed file or add agent will trigger autodiscovery on the system if it is successfully discovered.

Controlling autodiscovery with filters

Control autodiscovery with autodiscovery filters. Use the filters to add devices automatically or select the **Ask before adding new systems** option. When selected, this option places new devices in the Pending Devices list to be reviewed and then added manually.

You can further customize autodiscovery by making some filters automatic and some manual when adding systems to the topology. For example, automatically add routers and switches on subnets in a new phase of network expansion but continue to manually add all other devices. This use of autodiscovery filtering works best in networks that have consistent, well-defined naming conventions.

Autodiscovery filters are inclusive filters, so to prevent specific devices from being discovered, you must configure exclude filters. Use the ipExcludeList in the *discovery.conf* file to create these filters. Note that seed file and agent devices are not subject to the autodiscovery filters, but devices found by autodiscovery probing of the seed file or agent devices will be subject to the autodiscovery filters.

For more on filters see Chapter 6 of the *EMC Smarts IP Management Suite Discovery Guide*.

Automating manual discovery

Using the seed file or add agent without autodiscovery provides complete control over the discovery process while avoiding the additional autodiscovery probing. Though this method is usually a manual one, you can automate it: first, generate a seed file from an inventory system on a regular basis as the network changes. Then, import the file by invoking **sm_tpmgr --seed=<*seedfile*>** for the domain using **cron** or **sm_sched**. Other more sophisticated approaches can also be programmed.

Discovery and security

When planning discovery, consider the following network security-related features:

- Firewall Ports: If there is a firewall between any portions of the management infrastructure, certain TCP and UDP ports in the firewall must be opened for proper communications during discovery and for other EMC Smarts communications:
  - SNMP polls: 161
  - SNMP traps: 162
  - Broker: 426
  - License Manager: Port 1744
  - EMC Smarts Manager: one port each, which can be configured.
  - EMC Smarts adapters, including SNMP Trap Adapter (Receiver) and Syslog Adapter: See "Deploying Syslog Processing" on page 6-1 and "Designing Trap Processing" on page 6-1.

  Document the ports that are opened in the Deployment Build Guide.
- Use of access lists. If access lists are used, the IP addresses of servers running EMC Smarts products must be added to the access list of devices that will communicate with EMC Smarts. EMC Smarts must have full access to browse the MIBs of the devices.
- Use of SNMP versions and their respective security capabilities. The version of SNMP used to communicate with the network devices can provide dramatically different levels of security. With SNMP V1 or V2C, the security is provided through the use of SNMP community strings. To properly configure EMC Smarts, you must know the SNMP read community strings for all devices that will be managed.

  For communications to devices using SNMP V3, the requirements are much greater. Obtain values for these configuration parameters for each SNMP V3 device:
  - SNMP V3 user name
  - SNMP engine ID (optional)
  - Authentication protocol
  - Authentication password
  - Privacy protocol and password
  - Context name, if used

Discovery and certified device types

To ensure devices are certified in EMC Smarts, obtain a list of the manufacturers and models for all devices in the network. In some cases, it might be necessary to obtain a device MIB for certification. EMC Corporation certifies many devices, but some may be specialty devices for the particular organization (for example, private MIBs for SNMP agents in point of sales terminals) Document the types of devices to manage in the Deployment Build Guide.

Discovery and the Domain Name System

EMC Smarts relies on DNS as part of the automatic name resolution process for devices discovered in the topology. If DNS is not properly configured, the discovery process can be slowed considerably as EMC Smarts waits for DNS requests to time-out.

Both the forward *and* the reverse DNS lookup files must be complete and properly configured. Improper configuration of the reverse lookup pointer records is a common problem. As part of the discovery design, determine if the network administration will ensure the accuracy of the DNS configuration.

If you cannot rely on DNS, you must use the seed file to name devices in your network. This requires that you set the value of NameFormat in the name-param.conf file to "TM_USESEEDNAME". (See the *EMC Smarts IP Management Suite Discovery Guide* for details about name resolution.) Plan on creating a comprehensive seed file that includes all necessary names. In this configuration, if autodiscovery is enabled, seed names are not available for devices that are autodiscovered. The IP address will be used as the name for the autodiscovered device.

Advanced discovery post-processing

If needed, customized processing can be added to discovery post-processing using ASL rulesets. Additional post-processing steps are added by editing one of the supplied ASL rule sets. Each ASL rule set is configured in such a way that it is invoked during a specific phase of the discovery:

- Before a full discovery (*custom-start-fulldisc.asl*)
- Before a system is discovered (*custom-start-system.asl*)
- After a system is discovered (*custom-end-system.asl*)
- Before discovery post-processing (*custom-start-post.asl*)
- After discovery post-processing (*custom-end-post.asl*)

Typical applications of custom discovery post-processing include:

- A post-processing script could unmanage devices in IP address ranges or with specific name patterns before post-processing. If there are groups of IP Addresses that are not normally reachable, a list of IP ranges or some matching criteria could be used during post processing to ensure EMC Smarts will not unnecessarily ping these addresses. Devices can also be unmanaged individually from the Global Console.
- A post-processing script could rename devices after a system is discovered.

Discovery design checklist

Before discovering the network, the requirements in the following checklist must be completed. Each chapter in this guide includes a checklist. For ease of use, the checklists are all grouped together in "Design and Deployment Checklists" on page C-1.

Table 4-1 Discovery design checklist

Complete	Task	Description	Related documentation
Initial Discovery			
❑	Define a method for the initial topology discovery.	❑ Use a comprehensive seed file without autodiscovery. ❑ Use autodiscovery with an agent or seed file. Document the method in the Deployment Build Guide.	"Initial discovery" on page 4-2
Topology Maintenance and Subsequent Discovery			
❑	Define a schedule for Full Discovery.	Define a regular schedule for full discovery. Choose a time of relative inactivity. Document the schedule in the Deployment Build Guide. Include *crontab* or *sm_sched* control file entries if used.	"Topology maintenance and subsequent discovery" on page 4-4
❑	Define a schedule for Pending Discovery.	Define a regular schedule for pending discovery. Choose a time of relative inactivity. Document the schedule in the Deployment Build Guide. Include *crontab* or *sm_sched* control file entries if these utilities are used.	"Topology maintenance and subsequent discovery" on page 4-4
❑	Determine if Autodiscovery is appropriate.	Document choice in the Deployment Build Guide.	"Topology maintenance and subsequent discovery" on page 4-4
❑	Choose a method for adding devices to the topology.	❑ Agent without autodiscovery. ❑ Seed file without autodiscovery. ❑ Use autodiscovery with an agent or seed file. Document choice in the Deployment Build Guide.	"Adding new systems to an existing topology" on page 4-5
❑	Prepare Seed File or Choose Agent.	If a seed file will be used to add devices to the topology, obtain a list of devices with names or IP addresses. Document how to obtain the list or the location of the list in the Deployment Build Guide. If an agent will be used instead, document the IP address or name of the agent.	"Topology maintenance and subsequent discovery" on page 4-4
❑	Define Autodiscovery Filters.	If autodiscovery is enabled, configure autodiscovery filters. These are inclusive filters that add devices to the topology. Document the autodiscovery filter criteria in the Deployment Build Guide.	"Controlling autodiscovery with filters" on page 4-5
❑	Define an Exclude Filter.	To exclude specific devices, use the exclude filter in the *discovery.conf* file. This simplifies creation of the autodiscovery filters. Document exclude filter entries in the Deployment Build Guide.	"Controlling autodiscovery with filters" on page 4-5
❑	Obtain SNMP security parameters per device.	Domain Managers use SNMP to poll the device agents. In order to do this, the Domain Manager needs the appropriate security information for the SNMP version: V1 and V2C use READ community strings for every device that will be managed; V3 uses the user name, SNMP engine ID, authentication protocol and password, privacy protocol and password and the context name. These parameters will be needed during discovery. Document in the Deployment Build Guide if permitted.	"Discovery and security" on page 4-6

Table 4-1 Discovery design checklist (continued)

Complete	Task	Description	Related documentation
❑	Open necessary firewall ports.	If there is a firewall between any portions of the management infrastructure, certain TCP and UDP ports in the firewall must be opened for proper communications: • SNMP polls: 161 • SNMP traps: 162 • Broker: 426 • License Manager: 1744 • Domain Managers (1 per manager): configurable • EMC Smarts adapters, including SNMP Trap Adapter (Receiver) and Syslog Adapter: configurable Document the ports that are opened in the Deployment Build Guide.	"Discovery and security" on page 4-6
❑	Provide access to network devices to manage.	For each device that EMC Smarts will monitor, the device's access list must include the IP address of the hosts where EMC Smarts Managers are installed. EMC Smarts must have full access to browse the MIBs of the devices. Document in the Deployment Build Guide.	"Discovery and security" on page 4-6
❑	Ensure DNS is properly configured.	For EMC Smarts to name devices in its topology correctly, the DNS needs to be clean (proper forward and reverse lookup). If DNS is not used, use of an */etc/hosts* file or not doing any name resolution at all can be considered.	"Discovery and the Domain Name System" on page 4-7
❑	Determine if Discovery post processing is required.	Determine if discovery post processing using ASL rulesets will be used. Document in the Deployment Build Guide.	"Advanced discovery post-processing" on page 4-7
❑	List unreachable IP Addresses	If there are groups of IP addresses that are NOT normally reachable, assemble a list of IP ranges or some matching criteria so EMC Smarts will not unnecessarily ping these addresses. Document these addresses in the Deployment Build Guide.	"Advanced discovery post-processing" on page 4-7

5

Designing Polling and Thresholds

Organizations reflect their business priorities in their network organization. Certain parts of the infrastructure are more important and must be monitored more closely to ensure availability and performance. Critical systems may need to be polled more frequently to uncover problems sooner. This is reflected in polling and threshold groups.

Settings and groupings are determined by assessing the operational requirements for the delivery of timely analysis data and the impact of polling on network infrastructure components. Management policies are set by creating physical or logical groups of systems, ports, interfaces, etcetera, and specifying management parameters for each group.

The Availability Manager polls all systems while the Performance Manager polls switches and routers.

This chapter describes the polling and thresholds design considerations. For a complete discussion of polling and threshold groups, parameters, and default values, see the Groups and Setting chapters of the *EMC Smarts IP Availability Manager User's Guide* and/or *EMC Smarts IP Management Suite Configuration Guide*.

- Polling and polling groups 5-2
- Threshold values and threshold groups 5-3
- Polling and threshold checklist 5-4

Polling and polling groups

When creating polling groups, consider device types and roles such as core devices and edge devices. Consider functional roles such as routers supporting particular organizations or departments. For example, a Service Provider could provide various support levels that include different polling cycles. The Service Provider could charge more for more frequent polling assuming that it will result in an earlier resolution of root cause analysis.

Before modifying any polling groups, determine if the organization enforces any limitations on polling frequency. By default, devices are polled every 4 minutes. Note that some limitations may be based on CPU utilization limits or router traffic limits that no longer apply in the current network. Be sure that you understand how the limitations were determined and, if necessary, suggest that network administrators modify the limits.

Before modifying polling parameters for existing polling groups, ensure the changes are appropriate for all members in the polling group. If the proposed polling change affects only a limited number of devices in the group, and the current polling adequately supports most devices in the group, then consider creating a new polling group for the devices with the unique polling requirements.

Whenever new groups are created with more specific matching criteria than the existing groups, pay particular attention to the polling group priority. More specific groups should always be higher in priority than similar less specific groups because the first match is always used: a device is compared against each polling group's matching criteria and when a set of criteria corresponds, that group's polling parameters are used. No further comparisons are performed for the device.

The polling time-out specifies the amount of time the system waits for a response to a poll. The default to 700 milliseconds. Note that whenever you increase the polling time-out, pay attention to the effect on the cumulative value for the polling time-out. On each successive retry, the polling time-out is doubled. The cumulative polling time-out should always be less than the polling interval or the device could be polled excessively. Consider reducing the number of retries when increasing the polling time-out.

Considering network latency

Polling groups are especially important in multi-site environments managed from a single IP Availability Manager or IP Performance Manager. You must consider connection speeds and latency that may differ based on network location.

Refer to your network diagram to locate WAN links that EMC Smarts polls must cross to reach devices. Based on the latency of the links, you may have to adjust the polling cycles. These types of modifications are difficult to perform accurately before the actual deployment because the traffic load across the links that EMC Smarts polls may not be readily available. Expect to revisit the polling settings during the validation and tuning phase.

If latency becomes an issue that cannot be overcome with realistic polling groups, you must reconsider your design. If the network includes many lower speed links, it may be necessary to relocate one or more Domain Managers physically or even increase the number of Domain Managers to reduce the latency.

Threshold values and threshold groups

Threshold groups are organized in the Domain Manager into four main groups:

- Interface groups. Parameters can be set to monitor port performance, port flapping, backup support, and dial-on-demand support.
- Port groups - Truck ports. Trunk ports connect to other ports. Parameters can bet set to monitor port performance and port flapping.
- Port groups - Access ports. Access ports are all of those that are not trunk ports. Parameters can be set to monitor port performance and port flapping.
- System resource groups. Parameters can be set to monitor connectivity, environmental devices such as power supplies, fans, voltage, etcetera.

Judicious choice for thresholds can allow proactive rather than reactive network management, but poor choices can result in hundreds of inappropriate notifications.

It is difficult to adjust performance thresholds before gaining experience with the deployment and the network components being monitored.

Analyze failures of network components to determine if performance degradation was a precursor to failures. Based on the analysis, adjust the thresholds accordingly. This analysis should become a constant ongoing process with constant adjustment as performance notifications and failures occur.

In many cases, threshold values may be adjusted in response to inappropriate notifications.

Notifications from backup interface support thresholds and dial-on-demand interface support thresholds are some of the more common issues as these threshold groups are particularly unique to an individual network and its administrators.

Polling and threshold checklist

Each chapter in this guide includes a checklist. For ease of use, the checklists are all grouped together in "Design and Deployment Checklists" on page C-1.

Table 5-1 Polling and threshold checklist

Complete	Task	Description	Related documentation
❑	Determine polling group requirements.	Design polling groups based on importance of network device performance both to the network and to the various parts of the organization. Also consider network latency to determine if changes are needed. Document choices in the Deployment Build Guide.	"Polling and polling groups" on page 5-2
❑	Set polling parameters for each polling group.	Set polling parameters based on importance of network device performance. Additional modifications may be necessary if polling does not present an accurate picture of network availability during validation. Document new polling parameters in the Deployment Build Guide.	"Polling and polling groups" on page 5-2
❑	Determine threshold group requirements.	Design threshold groups based on importance of network device performance both to the network and to the various parts of the organization. Document choices in the Deployment Build Guide.	"Threshold values and threshold groups" on page 5-3
❑	Set threshold parameters for each threshold group.	Set threshold parameters based on the expected effect of degraded performance on network operations. Additional modifications may be necessary during validation and as the organization gains experience with the performance indicators. Document new threshold parameters in the Deployment Build Guide.	"Threshold values and threshold groups" on page 5-3

6

Deploying Syslog Processing

The Syslog Adapter reads the contents of any system log file and generates notifications based on the file contents. This chapter describes the choices that must be made when designing syslog processing. Complete configuration details are included in the *EMC Smarts Service Assurance Manager Adapter Platform User's Guide*.

- Syslog processing applications 6-2
- Syslog processing checklist (optional) 6-4

Syslog processing applications

The Syslog Adapter can read any text file in the proper format and parse the file to generate notifications for the Service Assurance Manager. Typical products that use the Syslog Adapter to monitor security violations at specific servers or monitor routing protocols are Border Gateway Protocol (BGP) and Open Shortest Path First (OSPF) protocol.

For example, the Syslog Adapter can listen to a system log (syslog) file that represents the combined syslog files for a group of similar routers. Then, when specified messages such as BGP adjacency changes are written in the syslog file, the Syslog Adapter generates notifications to the Service Assurance Manager via the Service Assurance Manager Platform.

The syslog messages that generate notifications and the corresponding notification's attributes are defined in the files *my_hook_syslog.asl* and *syslog_mgr.asl* in **BASEDIR**/*smarts/rules/icoi-syslog*.

Creating the syslog file

Consider the following questions about syslog processing before moving to EMC Smarts-specific deployment design:

- How will the syslog file be created? If the syslogs for a number of systems must be monitored, a typical application has the systems send the syslog messages to a UNIX host running a syslog daemon as a receiver for systems' messages. The syslog daemon then writes the messages to a combined syslog file. Applications other than syslog daemon can create files that the Syslog Adapter can read, such as the Event Log file for CiscoWorks on a Windows host.
- Where will the syslog file be located? The Syslog Adapter must be able to access the file. The monitored devices must be able to send messages to the system where the file will be compiled.

Processing the syslog file

In normal processing, the Syslog Adapter will tail the contents of a syslog file. When tailing a file, the adapter processes only new messages added to the file. Tailing provides constant monitoring of the syslog file while the adapter is running. If tailing is disabled, the Syslog Adapter parses and processes the file once.

Determine which messages in the syslog file are important in the deployment. Network administrators can explain any practices in their network that result in syslog messages and can recommend which syslog messages are appropriate for processing. Note that by default, only messages related to devices in the topology are used by the Syslog Adapter: all other messages are ignored. This behavior can be reconfigured, but processing other syslog messages may add to processing time significantly.

To expand syslog processing, choose additional messages that will generate notifications and determine what information the notifications will contain. Messages can be selected based on source and content to create notifications. With appropriate logic in the *my_hook_syslog.asl* file, information retrieved from the messages can be used to customize the notifications. You can use the Adapter Scripting Language (ASL) to modify *my_hook_syslog.asl* and specify the appropriate processing.

If you consider more extensive customization of notifications, remember that Syslog Adapter hook script processing is single-threaded: the more logic performed on each notification, the longer the processing time and potential bottleneck.

Revise your solution architecture diagram to respond to your syslog processing design. See "Syslog adapter added to the solution architecture diagram" on page 6-3 for a typical example.



Figure 6-1 Syslog adapter added to the solution architecture diagram

Syslog processing checklist (optional)

Each chapter in this guide includes a checklist. For ease of use, the checklists are all grouped together in "Design and Deployment Checklists" on page C-1.

Table 6-1 Syslog processing checklist

Complete	Task	Description	Related documentation
❑	Determine how to create the file for processing by the EMC Smarts Syslog Adapter.	A file must be created that the EMC Smarts Syslog Adapter can parse. Determine which devices will contribute messages to the file. Consistent layout of the messages in the file is required for Syslog Adapter processing. Include all details in the Deployment Build Guide.	"Creating the syslog file" on page 6-2
❑	Determine the location of the file that the Syslog Adapter will process.	The process that is creating the file must be able to receive messages from source applications and the created file must be accessible by the EMC Smarts Syslog Adapter. Include all details in the Deployment Build Guide, including location, host and path.	"Processing the syslog file" on page 6-2
❑	Choose the messages that are most important for processing.	Choose the messages that are most important for processing. Include all details in the Deployment Build Guide.	"Processing the syslog file" on page 6-2
❑	Determine the characteristics of the notifications that are generated.	For each message that generates a notification, determine the notification format. These characteristics will be used to develop the hook script for the EMC Smarts syslog processing deployment. Include all details in the Deployment Build Guide.	"Processing the syslog file" on page 6-2
❑	Add Syslog Processing to the Solution Architecture Diagram.	Add Syslog Processing to the Solution Architecture Diagram.	"Processing the syslog file" on page 6-2

7

Designing for Administration of EMC Smarts

This chapter describes how to categorize EMC Smarts users, choose their access rights, and then assign administrative or monitoring capabilities.

- Who are the Global Console users? 7-2
- Designing user profiles 7-4
- Designing notification lists 7-5
- Restricting console operations 7-6
- Designing consoles 7-7
- Planning for tools and tool deployment 7-8
- Planning for tools and tool deployment 7-8
- Administration Design Checklist 7-9

See the *EMC Smarts Service Assurance Manager Configuration Guide* for information about configuration and implementation of EMC Smarts access.

Who are the Global Console users?

When designing your deployment, determine who will use the Global Consoles and for what purposes. Initially, it is not necessary to define individual users—consider, instead, broad functional categories of users with similar needs and characteristics.

Name these categories and list their requirements in the deployment's Build Guide. Using specific position names as categories can make the process easier. For example, you might create one or more of these categories: network engineer, network administrator, network support specialist, technical specialist, NOC manager, NOC operator, LAN administrator, and IT management. Once you choose categories, determine the typical EMC Smarts-related duties that are performed by the personnel in these categories. You may find that as you list the duties, you may have to expand or combine certain categories. Once the categories are completely defined, list them in the Deployment Build Guide.

For example, Table 7-1 defines two typical EMC Smarts user categories.

Table 7-1 Typical EMC Smarts user categories

User Category	Description of Needs/Duties
Field Engineer	Administers and maintains local and wide area networks and related hardware. Monitors daily activity, enforces licensing agreements, and provides front line support, including both software and hardware support. • Needs monitoring access to all domains. • Needs to see all important traps, notifications, and network outages.
Local Area Network Administrator supporting a Customer of a Service Provider	Directs the daily operational availability of the hardware and software systems required to support facility operations. Directs and oversees scheduled testing and review of hardware and software to ensure potential problems are identified at the earliest point possible. Analyzes, evaluates and builds cost effective LAN solutions that leverage resources and technology to meet business requirements. Designs, creates and distributes user documentation relating to installation of software. • Needs monitoring access to the Domain Manager that supports the customer, but not access any other domains. • Needs to see all important traps, notifications, and network outages for the customer's EMC Smarts domain.

Users and security

When defining your functional groupings for users, you must consider EMC Smarts' security implementation. EMC Smarts is designed so that users fall into these access levels:

- All, a level where EMC Smarts users can access all Global Console functionality available for one or more Domain Managers, if their user profile permits it
- Monitor, a level where EMC Smarts users can access only Global Console monitoring functionality, not administrative functionality, at one or more Domain Managers, if their user profile permits it
- Ping, a level normally reserved for EMC Smarts processes, where processes will ping hosts where other EMC Smarts processes are installed to determine if the hosts are running.
- None, a level that specifically excludes access to the Global Console

Also consider which Domain Managers should be accessed by which users. You can define this access in the *serverConnect.conf* file on the servers where EMC Smarts is installed.

Password configurations

Determine how you will configure EMC Smarts passwords. You can do any or all of the following:

- Allow the host operating system to validate users. This method provides the highest level of security and is easy to manage because it relies on the security implementation that is already in place. There are two variations: any valid user can access one of the EMC Smarts levels (All, Monitor, Ping) or specific users can access a specific EMC Smarts security level. Defining specific access requires more maintenance because you must list the user names in EMC Smarts configuration files (*serverConnect.conf* and *clientConnect.conf*). See the *EMC Smarts System Administration Guide* for methods to configure and to secure access for these files.
- Specify unique EMC Smarts passwords for individual EMC Smarts users. Consider this method only when there are very few EMC Smarts users because it requires a high level of maintenance. Note that it is less secure than permitting the host to validate users.
- Specify a common EMC Smarts user name with a common password. This method is the least secure, but very easy to maintain.

Note that you can combine these methods, for example, you could restrict administration (All) capabilities to specific users validated by the operating system. In addition, you could provide Monitor level abilities to a general EMC Smarts user named “Monitor.”

Designing user profiles

The functional grouping of users and their requirements form the basis of EMC Smarts user profiles. These profiles combine access to notification lists, console operations, custom console layouts, and specific tools.

Create a profile for each category of user that you must support. Groups of EMC Smarts users with similar needs can then be assigned the same user profile.

If needed, you can further customize a generic user profile by copying it and then modifying it for more specific needs. For example, an administration user profile could be customized for less experienced administrators by restricting access to some administrative console operations and tools. Other possible user profiles could include regional or customer-specific consoles.

Designing notification lists

A Notification List determines the events that are forwarded to a user. Essentially, the list filters the notifications that are sent from the Service Assurance Manager and can be assigned to one or more users. The lists can be organized by:

- Business units
- Geographical regions
- Groups of resources

For example, a notification list can be defined to allow only notifications from the subnetworks devoted to a specific ISP customer to reach the Global Console of the customer's network administrator.

Restricting console operations

Most operations that can be performed at the Global Console can be individually enabled or disabled in the user profiles. When used with the security levels, restricting console operations can fine tune the abilities of users and further protect the EMC Smarts deployment.

For example, consider two users who have the Monitor security level that allows access to Global Console monitoring functionality, but not administrative functionality, for a Domain Manager. You can further restrict one of the viewers to see only the summary view and the IP Network map while restricting the other user to the notification log and the topology browser.

Designing consoles

The default EMC Smarts notification console is unfiltered, so EMC Smarts users may be overwhelmed by potentially enormous amounts of information. Properly configured filters can be used to customize notification consoles for groups of users. In addition, specific views can be automatically provided to match user needs.

Planning for tools and tool deployment

Many types of client and server tools can be designed and developed. Typically, tools will require different levels of programming skills based on their complexity. Plan to have personnel with the appropriate skills available.

In addition, determine which tools should be available to users through their user profiles.

Administration Design Checklist

Each chapter in this guide includes a checklist. For ease of use, the checklists are all grouped together in "Design and Deployment Checklists" on page C-1.

Table 7-2 Administration Design Checklist

Complete	Task	Description	Related documentation
❑	Define functional groups for users.	List the broad functional groups that users will belong to and then define the needs and duties of each group. Add this information to the Deployment Build Guide.	"Who are the Global Console users?" on page 7-2
❑	Define how you will implement user security.	Document in the Deployment Build Guide.	"Who are the Global Console users?" on page 7-2
❑	Define User Profiles.	Document in the Deployment Build Guide.	"Designing user profiles" on page 7-4
❑	Define notification lists.	Document in the Deployment Build Guide.	"Designing notification lists" on page 7-5
❑	Design console operations access.	Document in the Deployment Build Guide.	"Designing consoles" on page 7-7
❑	Design console layouts.	Document in the Deployment Build Guide.	"Designing consoles" on page 7-7
❑	Design client and server tools.	Document in the Deployment Build Guide.	"Planning for tools and tool deployment" on page 7-8
❑	Associate notification lists, console operations, console layouts, tools, and users with user profiles.	Document in the Deployment Build Guide.	"Administration Design Checklist" on page 7-9

8

Deploying EMC Smarts

This chapter describes important deployment considerations when installing EMC Smarts such as general installation procedures, licensing and security issues and deploying up trap processing.

This chapter should be considered as a supplement, not a replacement, for the *EMC Smarts IP Management Suite Installation Guide* and for the *EMC Smarts System Administration Guide*. Consult the *EMC Smarts IP Management Suite Installation Guide* for detailed installation procedures and consult the *EMC Smarts System Administration Guide* for detailed configuration procedures.

- General installation/deployment guidelines 8-2
- EMC Smarts licensing 8-3
- EMC Smarts installation 8-4
- Configure security 8-5
- Deploy trap processing 8-6
- Deploy EMC Smarts user configurations 8-7

General installation/deployment guidelines

Many organizations have strict rules for deploying enterprise-level software which may include deployment of staging areas whenever possible. A staging area is a copy of the installation.

Never install a deployment during a normal production shift; instead, choose a period of low utilization for installation. If a testbed is not available, create a staging area and make and test all changes in the staging area.

If possible, install the deployment in stages to reduce the size and complexity of each step in the overall process. Doing so makes debugging easier. Verify each stage of the deployment as described in the "Validating Your Deployment (Acceptance Testing)" on page 9-1.

Allow access to MIBs in network devices

The IP addresses of servers running EMC Smarts products must be added to the access list of devices that will communicate with EMC Smarts. EMC Smarts must be given full access to browse the MIBs that are listed in the *EMC Smarts IP Availability Manager User's Guide* and/or the *EMC Smarts IP Performance Manager User's Guide*.

EMC Smarts licensing

For permanent licensing, EMC Smarts products rely on FLEXnet licensing software and a license file provided by EMC Corporation after EMC Smarts software is purchased.

Note: Macrovision® Corporation changed the name of their license manager from FLEXlm to FLEXnet beginning with version 10. FLEXlm licenses are compatible with FLEXnet servers.

The FLEXnet software is installed with the EMC Smarts software, but a license file from EMC Corporation is also required. The license file must be installed using the *install_license* script as described in the *EMC Smarts System Administration Guide*. To generate the license, EMC Corporation needs the following information for the computer where the FLEXnet license server is running:

- The host ID (For Windows systems, use the Ethernet address)
- The hostname
- The operating system and version
- The number of systems that will be managed by EMC Smarts. Managed systems include routers, switches, hubs, virtual routers, security devices, hosts, servers, desktop or laptop PCs, workstations, probes, terminal servers, printers, IP phones, wireless access points (WAPs), and CSUs/DSUs

Note: An accurate count of systems ensures that the deployment can manage all the devices appropriately. The discovery process will halt if the number of discovered systems exceeds the licensed quantity.

Send this information to EMC Corporation as soon as possible after purchasing your EMC Smarts software to ensure that an appropriate license is available for the deployment.

When multiple license servers are used, each EMC Smarts Domain will use blocks of volume licenses from a specific license server. Therefore, the license blocks must be divided among multiple licenses and the appropriate license must be deployed to each license server. Then, Availability Managers that rely on a specific license server will have the appropriate volume licenses to support the systems that they must manage. If multiple license servers are used, ensure that EMC Corporation is aware of the configuration details.

EMC Smarts installation

The general rule is to install and configure items from the bottom of the hierarchy first, moving up to the Domain Manager. The suggested installation order is as follows:

1. EMC Smarts Broker (normally installed first, making it easier for the other components to connect to each other)
2. Underlying Global Managers (IP Availability Manager and IP Performance Manager)
3. EMC Smarts Adapters
4. Global Manager (Service Assurance Manager)
5. Global Console

After the Global Manager is installed and configured, you can install the components that use the Global Manager as a server, such as the Global Console.

To ease the troubleshooting of initial deployment, install more limited segments of the deployment first, such as an Availability Manager and then Global Manager. Always validate a segment before installing the next segment.

By default, all EMC Smarts products are installed as services and are started immediately after installation. During deployment, you should set the services to start manually until your installation and validation are complete.

Setting environment variables

Setting inappropriate values for environment variables is a common cause of post-installation problems. Detailed descriptions of the environment variables used by EMC Smarts, including methods for setting them, are described in detail in the *EMC Smarts System Administration Guide*.

Configure security

For initial validation, use the default administration user name and password (*admin* and *changeme*). Once you have validated your installation, change the default administration user name and password.

Use and guard your EMC Smarts secret phrase

EMC Smarts components are installed using a default secret phrase. This phrase can be used to encrypt EMC Smarts passwords used in authentication and to encrypt communications between EMC Smarts components.

EMC recommends that you take advantage of the added level of security provided through the secret phrase and its related security mechanisms. To do this, you must change the secret phrase using sm_rebond and make it consistent at all your installation sites. Due to the sensitive and vital nature of this secret phrase, store and guard the phrase as you would do with the root passwords of the most sensitive servers in your network.

Under certain circumstances, the loss of the secret phrase can force extensive reconfigurations and require reinstallations of all EMC Smarts components.

Deploy trap processing

In the recommended trap processing configuration, two SNMP Trap Adapter (Receiver) instances are invoked using different *trapd.conf* files. The *trapd.conf* for the "trap exploder" instance includes trap forwarding statements and indicates the port to use when listening for traps. In contrast, the *trapd.conf* for the other instance of the SNMP Trap Adapter does not include trap forwarding statements.

In the following procedure, *BASEDIR* is the location where the Service Assurance Management Suite is installed. Deploy the trap processing as follows:

1. Install the SNMP Trap Adapter (Receiver) as a service. The default configuration will use *BASEDIR/smarts/local/conf/icoi/trapd.conf*, the version of the *trapd.conf* file that is *not* configured to forward traps.
2. Create the startup configuration for the trap exploder instance of the SNMP Trap Adapter (Receiver) using **sm_service**. The trap exploder instance will use *BASEDIR/smarts/local/conf/trapd/trapd.conf*, the version of the *trapd.conf* file configured to forward traps. Typically, the startup configuration would look like this:

   ```
   ▼# BASEDIR/smarts/bin/sm_service install
      --startmode=runonce
      --description="SMARTS Trap Exploder" ic-trapd-exploder
      BASEDIR/smarts/bin/sm_trapd
      --name=EXPLODER-INCHARGE_OI --config=trapd
      --port=<number> --seed=seedfile --output --rules=default▲
   ```

   In this case, the value for *- -port* should be the port number where network devices send traps. For complete details on **sm_service**, see the *EMC Smarts System Administration Guide*.
3. Use **sm_edit** to configure the two different versions of the *trapd.conf:*
   - The SNMP Trap Adapter (Receiver) instance will use *BASEDIR/smarts/local/conf/icio/trapd.conf*. The file should not include trap forwarding statements.
   - The trap exploder instance will use *BASEDIR/smarts/local/conf/trapd/trapd.conf*. This file should include all trap forwarding statements. Traps that must be processed into notifications should be forwarded to *host:port* where the SNMP Trap Adapter (Receiver) instance listens for traps.
4. Use **sm_edit** to configure the *trap_mgr.conf* file for the SNMP Trap Adapter (Receiver) instance.
5. Start both the SNMP Trap Adapter (Receiver) and the trap exploder.

Deploy EMC Smarts user configurations

Deploying the EMC Smarts user configurations consists of two tasks:

- Configure access to EMC Smarts software by adding user names and passwords to the security files (*clientConnect.conf* and *serverConnect.conf*) as described in the *EMC Smarts System Administration Guide*. These files define the security level for each user, including user capabilities, servers that can be accessed, and passwords.
- Configure EMC Smarts users:
  - Deploy and configure tools. Server tools must be copied to the server where Service Assurance is installed and client tools must be copied to all systems where Global Consoles are installed.
  - Create console configurations and then save them on the server where Service Assurance is installed. Create each console configuration by opening the Global Console and then arranging the layout and customizing preferences. Save the customized console with an appropriate name in an **.iccon* file. Each console file can then made available to all users by copying it from *BASEDIR/smarts/local/consoles/<user>* to *BASEDIR/smarts/local/consoles*.
  - Define customized notification lists.
  - Create user profiles which associate individual users with access to specific tools, consoles, and a notification list.

9

Validating Your Deployment (Acceptance Testing)

This chapter describes how to ensure that your EMC Smarts deployment operates as intended.

- Validation techniques 9-2
- Initial validation 9-3
- Validating discovery 9-4
- Validating polling and events 9-5
- Validating trap processing 9-6
- Validating users and capabilities 9-7

Validation techniques

When validating, begin by dividing the deployment into manageable, logical segments. Ensure each of the segments function properly and then perform end-to-end testing. Check data flow and then check the accuracy of the data itself. Validate as much as possible before discovering the topology so that you reduce complexity.

Initial validation

To start validation, ensure that the Broker has access to all Domain Managers and that the Domain Manager processes are registered and running. Use brcontrol, as described in the *EMC Smarts System Administration Guide*, to list the EMC Smarts processes registered with the Broker and their status.

Validating discovery

Discover the topology of the network. If the deployment includes a single Availability Manager on a large or extra large platform, discover the network in limited portions: for example, discover groups of one thousand managed network devices. Monitor the discovery process and review the discovered topology after each discovery.

If the network includes Hot Standby Router Protocol (HSRP) groups or virtual routers, ensure they are discovered.

For Availability Manager and Performance Manager, validate the discovered topology by reviewing a segment of the network that is well known. Always confirm possible discovery errors: if an expected device does not appear in the topology, ping the device to ensure it is accessible. It is not unusual for discovery to find more devices than you expect; if this is the case, confirm that the devices exist.

Validating polling and events

To validate polling and thresholds, do the following using the Polling and Thresholds Console:

- Cause a failure by physically removing a cable in a discovered portion of the network. Check that the correct notification is received at the Global Console and that the topology map indicates the failure. Note that you may have to wait a few polling cycles to see the correct root-cause analysis.
- Reduce threshold settings to very low levels or zero. Typically, reducing port or interface performance settings works well for validating Performance Manager deployments. For Availability Manager deployments, reduce the RestartTrapThreshold connectivity setting for a router or switch and cause a warmstart on the corresponding type of system. Obviously, choose equipment and a time frame when you will not interrupt your network users. Verify that the correct notification is received at the Global Console indicating that the threshold was exceeded.

Validating that the levels that you chose for the thresholds are appropriate for your deployment is much more difficult. Start by using the defaults and adjust them upward or downward as you gain experience with the equipment. When failures occur, particularly failures with little or no warning, determine if the failures were preceded by symptoms that could have been detected by lower threshold values. Then adjust the thresholds appropriately. This evaluation should be performed after all failures.

Validating trap processing

Use sm_snmp to generate traps to the SNMP Trap Adapter (Receiver) functioning as the trap exploder.

Ensure that you send traps that cause notifications from the underlying analysis servers as well as from the SAM Adapter configuration. Create a trap for each trap processing statement in the trap_mgr.conf file. If ASL scripting is included in the trap processing, ensure the scripts function as intended.

Validating users and capabilities

To validate users and their capabilities, test each user profile. Create a temporary user for each of your user profiles. Log onto the Global Console as each user in turn and review the associated console capabilities including access to console operations, access to tools, configuration of the notification list, and layout of the console. Ensure the capabilities match your expectations for each user profile.

10

Tuning Your Deployment to Improve Performance

This chapter describes methods for identifying performance problems in a deployment and the steps that you can take to correct these problems.

- Performance tuning guidelines 10-2
- Assessing performance 10-3
- Reviewing EMC Smarts license metrics 10-4
- Reviewing EMC Smarts performance metrics 10-5
- Improving performance 10-10
- Other tuning issues 10-11

Performance tuning guidelines

Whether tuning of EMC Smarts software is required is directly related to the capabilities of the equipment where the EMC Smarts products are installed.

When a deployment is installed on equipment from the small and medium platform equipment tiers, defined in "Determine the required size of the EMC Smarts deployment" on page 3-3, tuning is usually not required. But performance should be monitored to ensure that there are no issues.

When the deployment is installed on equipment from the large or extra large platform equipment tiers, tuning for peak performance is almost always required. For these large and complex deployments, expect that you will require aid from the EMC Professional Services organization for the tuning process.

General tuning recommendations and details on network sizing are described in Table 10-1.

Table 10-1 Tuning recommendation for various sized EMC Smarts platforms

Platform Equipment Tier Size	Small	Medium	Large	Extra Large
Managed Ports and Interfaces	AM: 10K AM+PM: 5K	AM: 25K AM+PM: 15K	AM: 50K AM+PM: 25K	AM: 100K AM+PM: 50K
Hardware Configuration	2-4 CPUs (low-end) 1GB RAM	4 or more CPUs 2GB RAM	16 or more CPUs 4GB RAM	32 or more CPUs (high-end) 8GB RAM
Fine-tuning required?	Usually never	Possibly	Yes	Absolutely
Additional Sizing Comments	Add a second CPU for growth or if collocating a PM.		Use high-end CPUs for growth or if collocating a PM.	Consider splitting your topology into multiple domains.

AM = Availability Manager
PM = Performance Manager
AM+PM = Availability Manager and Performance Manager as a single process

Note: SNMP V3 authorization protocol, requires more computer resources than earlier version of SNMP. If you are using a large number of SNMP V3 devices or systems, you can expect slower performance than indicated here.

Regardless of the size of your deployment, always monitor its performance. If tuning is required, remember that tuning your deployment is an ongoing process that should be performed regularly. Network changes can potentially affect the performance of EMC Smarts: if you add, remove, or relocate network equipment, review the performance metrics.

Never waste resources by tuning a partial deployment—adjustments made during a partial deployment will usually be inappropriate for a complete deployment.

Assessing performance

Performance is assessed by checking the resources required by EMC Smarts processes against the operating system limits and by reviewing EMC Smarts performance metrics.

Checking resource requirements against operating system limits

The memory limit is an operating system limit that indicates the maximum process size permitted in memory. This limit may become an issue for very large topologies. The process size limits are shown in Table 10-2.

Table 10-2 Operating system limits on total process size in virtual memory

Operating System	Memory Limit
Solaris	3.75 GB
HP-UX	2 GB
Linux	3.8 GB
Windows	2 GB
See the appropriate EMC Smarts Installation Guide for specific operating system version and patch information. For more information on operating system limits, consult the operating system's documentation.	

To find the total size in memory for various EMC Smarts processes, use the following commands:

- For Linux or UNIX systems, use ps -l -p *proclist* where *proclist* is one or more IDs for the EMC Smarts processes. Review the total size of the processes in virtual memory (usually the SZ column).
- For Windows systems, open the Windows System Information tool and find the process size in the Software Environment, Running Tasks folder.

If an EMC Smarts process is approaching the memory limit, consult with EMC Professional Services for further evaluation.

Reviewing EMC Smarts license metrics

To determine if your deployment has sufficient licenses, Availability Manager and Performance Manager provide the sm_tpmgr utility. Use this utility with the following syntax to generate a complete list of performance metrics and deployment size information:

```
sm_tpmgr -s servername --sizes
```

For example, lines similar to the following would appear in the output:

```
Total System Volume License Checked Out:         150
Total Systems in Topology:               105
Remaining Blocks of System Licenses in License Server: 5
Maximum Number Of Systems: 1500
```

In this example, licenses for 1395 additional systems are available.

Reviewing EMC Smarts performance metrics

To assess performance, Availability Manager and Performance Manager provide the sm_tpmgr utility. Use this utility as follows to generate a complete list of performance metrics:

```
sm_tpmgr -s servername --show-dm-processes
```

The output of this command lists the duration of most tasks performed by the Domain Manager, including:

- Codebook tasks
- Discovery cycles, including post processing, reconfiguration, and saving the repository
- ICMP statistics
- SNMP statistics

Note that when collecting performance statistics for Availability Manager, useful statistics are only available after Availability Manager is in a steady state for two to three hours.

Codebook tasks

The codebook tasks are single-threaded and CPU-bound. These tasks occur each time discovery is completed, when reconfiguration occurs, or when topology changes are made. The tasks may also be triggered manually.

The codebook task durations are listed at the beginning of the output from the sm_tpmgr utility and will be similar to "The sm_tpmgr utility: codebooks tasks" on page 10-5.

```
     Last Consistency Phase One Update = 1
       Last Offline Consistency Update = 1
                Last Consistency Update = 0

     Last Offline Codebook Computation = 2
             Last Codebook Computation = 2

                       Last Correlation = 0

                         codebookRadius = -1
                   correlationInterval = 30
                     correlationRadius = 4

  Process                 Last Start Time                 Last End Time          Duration
Phase One Consistency  03-Nov-2003 19:15:00 EST   03-Nov-2003 19:15:01 EST          1
 Offline Consistency   03-Nov-2003 19:15:03 EST   03-Nov-2003 19:15:04 EST          1
        Consistency    03-Nov-2003 19:15:06 EST   03-Nov-2003 19:15:06 EST          0

  Offline NewMatrix    03-Nov-2003 19:15:06 EST   03-Nov-2003 19:15:08 EST          2
          NewMatrix    03-Nov-2003 19:15:06 EST   03-Nov-2003 19:15:08 EST          2

        Correlation    03-Nov-2003 19:15:18 EST   03-Nov-2003 19:15:18 EST          0
```

Do not use this value.

Refer to these codebook task values.

Figure 10-1 The sm_tpmgr utility: codebooks tasks

Look for these codebook task values to evaluate performance:

- Consistency: The time (seconds) required to recalculate aggregations.
- NewMatrix: The time (seconds) required to re-compute the codebook.
- Correlation: The time (seconds) required to perform the correlation

Table 10-3 indicates the acceptable and unacceptable values for the codebook tasks.

Table 10-3 Evaluating codebook task durations

Performance Category	Performance Assessment			How to Improve
	Good	Acceptable	Unacceptable	
Consistency	< 10 sec	10-30 sec	> 30 sec	Use a faster CPU or split the topology.
NewMatrix	< 10 sec	10-30 sec	30 sec	Use a faster CPU or split the topology.
Correlation	0	0-20 sec	20 sec	Use a faster CPU or split the topology.

Duration of last discovery

When assessing discovery performance, do not use the numbers from an initial discovery process because it is not representative of typical discovery processes and is usually very costly. In addition, discovery durations vary dramatically based on the type of discovery processing:

- Discover all
- Discover pending
- Discover one device
- Discover many devices using a seed file

Each of these must be assessed separately. You can find the time and duration of the discoveries in the Topology Manager section of the output from the **sm_tpmgr** utility. It lists discovery times and durations similar to that shown in Figure 10-2, "The sm_tpmgr utility: discovery information in the log file."

```
Topology Manager:
                    durationOfLastProbe = 0 00:09:00   Length of Discovery
                   lastProbeFinishedAt = 03-Nov-2003 07:15:00 PM EST
                    lastProbeStartedAt = 03-Nov-2003 07:06:00 PM EST
             lastProbeStartedAt_pending = 03-Nov-2003 07:06:00 PM EST
                        numberOfAgents = 0
                     numberProbeThreads = 10
                        probeQueueSize = 0
```

Figure 10-2 The sm_tpmgr utility: discovery information in the log file

If a full discovery takes more than eight hours (including post-processing and reconfiguration), it may affect normal business operations. In an environment where full discovery can be scheduled over a weekend, eight hours or more might be acceptable, but between two and five hours is often required to avoid interfering with possible third shift work.

Note: In addition to the sm_tpmgr utility, the DiscoveryInProgress event could be used with ASL scripting to determine the length of the task.

Discovery post processing

The log files provide the most accurate source of information for discovery post processing. In the BASEDIR/smarts/local/logs directory of the Availability

Manager, Performance Manager, or Discovery Manager, there is a *server_name.log* file where *server_name* is name of the Domain Manager.

Search the log file for the most recent "Started basic post-processing" statement and the most recent "Finished partitioning" statement. The duration between the times of these statements is the length of time required for discovery post-processing.

Do not use the discovery postprocessing duration when a large number of new devices are discovered and added to the topology because the post-processing required to create a new topology is much higher than under normal circumstances.

Additionally, in large or very meshed topologies, the greatest cost of partial discoveries, even of a single device, is in discovery post-processing. If discovery post processing is taking too long, review any custom post processing. If it is poorly designed or implemented, it should be reconsidered before splitting the topology or adding CPUs.

Table 10-4 Evaluating codebook task durations

Performance Category	Performance Assessment			How to Improve
	Good	Acceptable	Unacceptable	
Post-Processing	< 5 min	5-10 min	> 10 min	Use a faster CPU or split topology. Review any custom post-processing.

Reconfiguration and saving the repository

Reconfiguration is single-threaded and CPU-bound. Saving the repository is also single-threaded, but is mostly I/O-bound. These task durations appear in the output from the sm_tpmgr utility and will be similar to "The sm_tpmgr utility: reconfiguration and repository save tasks" on page 10-7.



Figure 10-3 The sm_tpmgr utility: reconfiguration and repository save tasks

Table 10-5 indicates the acceptable and unacceptable values for these tasks.

Table 10-5 Evaluating reconfiguration and repository save tasks durations

Performance Category	Performance Assessment			How to Improve
	Good	Acceptable	Unacceptable	
Reconfigure	< 5 min	5-8 min	> 8 min	Use a faster CPU or split the topology.
Repository Save	< 1 min	1-3 min	> 3 min	Use a faster CPU or split the topology.

ICMP processing statistics

ICMP processing statistics appear in the output of the sm_tpmgr utility and will be similar to "The sm_tpmgr utility: ICMP statistics" on page 10-8.

Look for these statistics to evaluate performance:

- avg_get_time, max_get_time, min_get_time: The average, maximum, and minimum duration (seconds) of a ping cycle.
- avg_late_polling: The amount of time (seconds) that the ICMP pinger is falling behind. If the value is negative, the pinger is ahead of schedule.

```
Pinger Accessor Interface:
                      avg_late_polling = -0.01840490847826
                           bytesPerPing = 64
        gets_causing_request_percentage = 0
              gets_from_cache_percentage = 100
                       icmpNumberOfPolls = 90
                    icmpNumberOfResponse = 90
                      icmpPollerTimeSkew = -0.0184049
                           icmpStartTime = 31-Oct-2003 10:10:44 EST
                      num_other_failures = 0
                             num_threads = 1
                            num_timeouts = 0
                          operation_size = 1
                periodic_gets_per_second = 0
               total_active_poll_actions = 0
                         total_get_nexts = 0
                   total_gets_from_cache = 162
     total_instrumentation_get_requests = 0
                    total_on_demand_gets = 0
                     total_periodic_gets = 0
                      total_poll_actions = 0
                        total_repos_gets = 162
```

Average Late Polling

Figure 10-4 The sm_tpmgr utility: ICMP statistics

Table 10-6 indicates the acceptable and unacceptable values for these statistics.

Table 10-6 Evaluating ICMP statistics

ICMP Performance Category	Performance Assessment			How to Improve
	Good	Acceptable	Needs Improvement	
avg. get time, max get time, min get time	<1 msec		>1 msec	Increase polling interval or split the topology.
avg. late polling	Negative	0-0.5 sec	> 0.5 sec	Increase polling interval or split the topology.

SNMP processing statistics

SNMP processing statistics appear in the output of the sm_tpmgr utility and will be similar to "The sm_tpmgr utility: SNMP statistics" on page 10-9. Look for these statistics to evaluate performance:

- avg_get_time: The average time (seconds) that an SNMP get cycle takes on a per-thread basis (that is, the round-trip delay). The value varies based on the network configuration.
- avg_late_polling: The amount of time (seconds) that the SNMP poller is falling behind.
- avg_request_size: The average number of variable bindings in a get request. Higher values indicate more efficient polling.
- num_threads: The number of polling threads.
- periodic_gets_per_second: The actual throughput of the poller across all threads.

```
SNMP Accessor Interface:
            avg_active_processing_time = 0.0973929232857143
                          avg_get_time = 0.0412507578666667
                     avg_late_polling = 0.188394353857143
                         avg_lock_wait = 0.000147875857142857
                     avg_request_size = 14
       gets_causing_request_percentage = 0
            gets_from_cache_percentage = 100
            max_active_processing_time = 0.369101803
                          max_get_time = 0.091588776
                     max_late_polling = 0.805186757
                         max_lock_wait = 0.000490496
                          min_get_time = 0.001716586
                   num_other_failures = 0
                          num_threads = 10
                          num_timeouts = 0
             on_demand_gets_percentage = 0
                        operation_size = 19
             periodic_gets_per_second = 5.13920718548761E-05
             periodic_gets_percentage = 100
             total_active_poll_actions = 7
                       total_get_nexts = 0
                   total_gets_from_cache = 744
    total_instrumentation_get_requests = 15
                  total_on_demand_gets = 0
                   total_periodic_gets = 15
                    total_poll_actions = 7
                      total_repos_gets = 744
```

Figure 10-5 The sm_tpmgr utility: SNMP statistics

Improving performance

If the performance metrics indicate a performance degradation, try the following tactics to improve performance:

- Split topology into multiple domains. Currently, EMC Professional Services can aid in efficiently splitting a network topology across multiple Domain Managers.
- Improve the capabilities of the equipment where the EMC Smarts process is installed: Use a faster CPU or reinstall EMC Smarts on more capable equipment. Before resorting to this hardware upgrade, consider the previous option.

Other tuning issues

Adjust performance thresholds to reduce inappropriate alarms

After gaining experience with the EMC Smarts deployment, adjusting performance thresholds may reduce the number of inappropriate alarms. Devices may trigger alarms during normal operation because performance thresholds set inappropriately low. Review all performance-related alarms that were triggered by conditions that did not represent actual or potential failures and adjust the threshold to avoid repetition of the alarm.

Conversely, failures of some devices may be preceded by performance degradation that does not trigger an alarm. Once again, review any failures that may be preceded by degraded performance and adjust the thresholds to detect these changes appropriately.

Use batching to improve trap processing performance

In a deployment where a high frequency of traps is expected, plan on using the batching capability of the EMC Smarts Service Assurance Manager Adapter Platform to improve performance of the clients that process the notifications. The *BATCH_NOTIFY_INTERVAL* in the *trap_mgr.conf* configuration file determines the length of the interval between sending batches of notifications based on traps. It may be necessary to tune this value under the typical trap load, so plan on monitoring the client performance and adjusting this value.

Global Console performance

To ensure appropriate performance at the Global Console, limit the total number of Global Console users attached to the same Service Assurance Manager to 50. This is typically a safe limit, but more users are possible. Check processor utilization to ensure it is not too high before adding additional users.

A typical Global Console tuning task is to optimize notification list filters for the operators. The operator should not see more (or less) than needed. Avoid needless processing whenever possible.

A

Using Discovery Manager to Obtain Network Information

This appendix explains how to deploy the EMC Smarts IP Discovery Manager to obtain the network information needed when designing a deployment.

- Deploying temporarily A-2
- Deploying permanently A-3
- Discovery process A-4
- Exporting topology A-5

Deploying temporarily

EMC Corporation offers a traveling license for the EMC Smarts IP Discovery Manager which allows you to deploy the IP Discovery Manager on a laptop. The laptop can then be connected to an organization's network and be used to discover IT objects in the networks.

You may also use the travelling license to temporarily install the IP Discovery Manager on a platform in an organizations's network. This is more cumbersome as it requires installation and configuration each time it is used in a new network. In addition, you must uninstall the IP Discovery Manager before moving on to another network.

Typically, the IP Discovery Manager is used to show due diligence in a temporary deployment.

Deploying permanently

The IP Discovery Manager can also be deployed in a network permanently using a static license. This deployment is beyond the scope of this Appendix.

Discovery process

As described in Chapter 4, "Planning for Discovery,", IP Discovery Manager's autodiscovery feature can be used to discover the network. However, this uses more network resources that any other method. Therefore, this method still requires that you carefully plan and design a discovery process as explained in "Planning for Discovery" on page 4-1.

Exporting topology

The topology can be retrieved with dmctl or by using an ASL script. In general, users need to develop their own tool to export the topology in the format they want. Users can also use sm_topodump to dump a subset of the topology or they can export the quantities of various network components using the sm_tpmgr.

Using sm_tpmgr

After discovering the network, use sm_tpmgr --sizes to retrieve data on the topology. sm_tpmgr --sizes breaks down topology by major classes of objects (Routers, Switches, Interfaces, Ports, Cables, etc.) and identifies what is instrumented (that is, monitored) and for which product (IP Availability Manager, IP Performance Manager or both).

For example, from ***BASEDIR****/smarts/bin*, enter:

```
% ./sm_tpmgr -s INCHARGE-AM --sizes
```

Typically the output from *sm_tpmgr -- sizes* is as follows:

```
Total System Volume License Checked Out:         150
Total Systems in Topology:               105
Remaining Blocks of System Licenses in License Server: 5
Maximum Number Of Systems: 1500

Number of Systems [Instrumented for Connectivity/Performance] Total
   Number of Systems: 105 [105/0]
        Number of Node: 10 [10/0]
        Number of Probe: 1 [1/0]
        Number of Switch: 11 [11/0]
        Number of Host: 78 [78/0]
        Number of Router: 4 [4/0]
        Number of Hub: 1 [1/0]

Total Number of IPs: 172 [170/0]
Total Number of Ports: 427 [91/0]
        Number on Switch: 427 [91/0]

Total Number of Interfaces: 160 [157/0]
        Number on Node: 11 [11/0]
        Number on Probe: 10 [10/0]
        Number on Switch: 26 [25/0]
        Number on Host: 94 [94/0]
        Number on Router: 18 [16/0]
        Number on Hub: 1 [1/0]

Total Number of Links: 76
        Number of NetworkConnections: 1
        Number of Cables: 71
        Number of TrunkCables: 4

Total Number of MACs: 648
Total Number of STPNodes: 43
```

B

Managing Overlapping IP Networks With EMC Smarts

EMC Smarts enables service providers who offer managed IP network services the ability to centrally manage the private IP networks of customers who deploy identically-numbered IP address spaces.

- Using IP management domains to manage overlapping IP networks.....B-2
- Consolidating management domain information.....B-7

To provide this capability, EMC Smarts must be deployed in an environment where the following criteria are met:

- The system hosting the EMC Smarts product supports virtual IP interfaces.
- A policy-based router between the system hosting EMC Smarts and the customer networks that supports policy routing based on a packet's source address.

Note: As an alternate to the method described in this chapter, EMC Smarts IP Management Suite version 7.0 has implemented an VRF-IP tagging capability that allows discovery of overlapping IP address as are often found in an MPLS VPN environment. See the *EMC Smarts IP Management Suite Discovery Guide* for details.

Using IP management domains to manage overlapping IP networks

To use EMC Smarts to manage overlapping IP networks, you must establish separate IP Management Domains. An IP Management Domain is a set of IP networks that do not contain overlapping addresses.

Consider two customers, A and B, that both use the private IP network number 10.0.0.0/8 (subnet 255.0.0.0). In this case, Customer A is one management domain, Domain A, and Customer B is a second domain, Domain B. The combination of A and B cannot be a domain because they both use the 10.0.0.0/8 network number. As the administrator of the EMC Smarts product, you map each customer network to a different management domain when configuring the EMC Smarts products.

EMC Smarts assumes that IP addresses within a domain are unique (except for cases such as Hot Standby Routing Protocols where it is possible for two devices to share the same virtual IP address so long as only one device actually uses it at any one time). IP addresses do not have to be unique across all of the management domains.

Topology and event information for the systems and networks in each IP Management Domain are kept isolated by running separate EMC Smarts products on the same host system. Because the topology and event information for each domain is separated, EMC Smarts does not display phantom connectivity between devices in different domains, even if they share the same IP network numbers. You can manage each domain from the same Administration Console, but the settings for each domain are separately configurable.

An IP Management Domain provides the structure into which the different customer networks are placed. But to actually discover the networks and then to subsequently poll them and process traps requires specialized support provided by the EMC Smarts product, as well as configuration support from the host operating system.

To manage networks with overlapping IP addresses, you need to complete the following steps.

- Configure the EMC Smarts products, one per IP Management Domain, by binding a virtual IP interface to each Domain Manager.
- Configure a policy-based router, or use source routes, to route packets sent from the EMC Smarts products to the appropriate IP Management Domain properly.
- Configure the devices in each IP Management Domain to send SNMP traps to the virtual IP address for which the respective EMC Smarts product is listening for traps.

Using virtual IP interfaces to direct management traffic

You bind an EMC Smarts product, and thus an IP Management Domain, to a virtual IP interface on the host operating system. Consider again the example from the previous section that defined two domains, Domain A and Domain B. To discover and poll these networks, the administrator must configure a separate virtual IP interface—Virtual Interface A and Virtual Interface B in this example—for each domain. Then, as part of the definition of the domain within the EMC Smarts product, you bind a domain to a particular virtual interface. In this case, Domain A may be bound to Virtual A and Domain B to Virtual B.

The domain-to-interface binding has the following effect on the requests coming into and going out of the EMC Smarts product:

- All SNMP and ICMP requests from Domain Manager A specify Virtual A as the source address. Requests from Domain Manager B specify Virtual B as the source address. Because of this, two packets destined for identically addressed devices in Domain A and Domain B are differentiated by the packet source address, as shown in "Path of packets delivered to EMC Smarts applications" on page B-3.



Figure B-1 Path of packets delivered to EMC Smarts applications

- An SNMP response or trap whose destination address is Virtual A is interpreted in the context of Domain A. The response or trap whose destination is Virtual B is interpreted in the context of Domain B. Incoming traps from identically addressed devices in Domain A and Domain B are distinguishable by the packet destination address as shown in "Path of packets delivered to EMC Smarts" on page B-3.



Figure B-2 Path of packets delivered to EMC Smarts

Proper trap processing support requires that devices in Customer A's network are configured to send traps to Virtual A. Devices in Customer B's network must be

configured to send traps to Virtual B. Because responses are sent to the source address of the query that initiated them, the rules for sources described previously ensure that responses are automatically sent to the right destination.

Binding an EMC Smarts application to a virtual interface

First, create one or more virtual IP addresses on the host where the EMC Smarts products are running. After you have created the virtual interfaces, you can bind each EMC Smarts product to a virtual interface. The Domain Manager is bound to the IP address at startup. Because of this, you will need to modify the startup script for each Domain Manager.

The sm_server command, which is used to start a Domain Manager, provides a special option, `--useif`, for binding the Domain Manager to a particular interface. The following example illustrates this option by binding the Domain Manager INCHARGE_A to the IP address 192.168.1.2.

```
# BASEDIR/smarts/bin/sm_server -n INCHARGE_A --useif=192.168.1.2
```

Routing management traffic to and from the management domains

In addition to binding each EMC Smarts product to a virtual IP interface, you need to configure a router to route packets to and from each EMC Smarts product to the appropriate managed domain. You can use one of two methods to route management traffic.

- A policy-based router that routes traffic from each virtual interface through interfaces that are only connected to each managed domain.
- Source routing where routing instructions are added to the packets leaving each EMC Smarts product.

Using a policy-based router to route management traffic

A policy-based router using methods such as Multiprotocol Label Switching (MPLS) or Border Gateway Protocol (BGP) is required to properly route the packets coming out of the EMC Smarts products. "Packets routed through a policy-based router" on page B-5 shows this example, where such a router must have two interfaces, Interface A and Interface B, through which there is unambiguous connectivity to networks of Customers A and B, respectively.

A policy route must be defined such that a packet, originating from Domain Manager A whose source address is Virtual A, is forwarded out of Interface A. Similarly, a second route must be defined such that a packet, whose source address is Virtual B, is forwarded out of Interface B. Packets returning from the customer networks through Interface A and Interface B are routed back to the appropriate interface and EMC Smarts product using standard IP routing.



Figure B-3 Packets routed through a policy-based router

Using source routes to route management traffic

You can also use either loose or strict source routes to route management traffic. The advantage of source routing is that it you do not have to configure the routers. However, not all routers support source routes. Figure B-4, "Packets containing routing instructions," shows an example where routing instructions are added to the packets leaving the EMC Smarts product. Each packet arrives at a router with instructions about the packet's next destination. When the instructions are exhausted, the packet will be in a location where standard routing can complete the packet's delivery.



Figure B-4 Packets containing routing instructions

Configuring devices to send SNMP traps

The final step to managing networks of overlapping IP addresses is to configure the devices in each network to forward SNMP traps to the appropriate EMC Smarts

product. Devices should be configured to send SNMP traps to the IP address of the appropriate virtual interface used by the Domain Manager. In our example, devices in Domain A are configured to send SNMP traps to the virtual interface used by Domain Manager A.

The configuration of EMC Smarts' trap processing is described in *EMC Smarts Service Assurance Manager Adapter Platform User's Guide*. EMC Smarts' trap processor runs automatically as a process within the Domain Manager. If the configuration for both trap processors is the same, they can share the same *trapd.conf* configuration file. Because each Domain Manager will listen for traps on different IP addresses, both Domain Managers can listen on the same port.

Consolidating management domain information

Once the IP Management Domains have been defined, and the virtual IP interfaces and policy routes established, each EMC Smarts product separately monitors and correlates the information for each domain.

If you want to consolidate the topology and event information to a single point of reference, use Service Assurance Manager. Service Assurance distinguishes between topology elements with the same IP address, providing separate notifications and a distinct topological representation for each element.

C

Design and Deployment Checklists

This Appendix gather the checklists that appear throughout this guide into a single, easily accessible location.

- Before you begin checklist C-2
- Architectural information checklist C-3
- Solution architecture diagram checklist C-5
- Discovery design checklist C-6
- Polling and threshold checklist C-8
- Syslog processing checklist (optional) C-9

Before you begin checklist

Before you begin a deployment, you must meet the requirements described in the following checklist.

Table C-1 Before you begin checklist

Complete	Requirement	Description
❑	Possess an understanding of the EMC Smarts architecture and capabilities.	At a minimum, you must understand the concepts and EMC Smarts architecture described in the following documents: • *EMC Smarts IP Availability Manager User's Guide* • *EMC Smarts IP Management Suite Discovery Guide* • *EMC Smarts IP Management Suite Configuration Guide* • *EMC Smarts IP Performance Manager User's Guide* • *EMC Smarts Service Assurance Manager Configuration Guide* • *EMC Smarts Service Assurance Manager Introduction* • *EMC Smarts Service Assurance Manager Adapter Platform User's Guide* • *EMC Smarts System Administration Guide* • *EMC Smarts IP Management Suite Installation Guide* that accompanied your software product suite
		To improve your understanding, attend EMC Smarts training courses offered by EMC Education Services & Development. Typically, deployment requires the knowledge equivalent to what is provided in the training courses on: • EMC Smarts IP Management Suite • EMC Smarts Service Assurance Manager • EMC Smarts Service Assurance Manager Adapter Platform
❑	Obtain contact information for the deployment team.	The contact list should include titles, responsibilities, and contact methods for all team members.
❑	Get nondisclosure requirements and negotiate an agreement.	Be aware of the requirements of the non-disclosure agreements that are in place for the EMC Smarts deployment.
❑	Develop schedules and set milestones for early deliverable.	Scheduling a software deployment varies based on the size and scope of the deployment and the organization's requirements. Typical milestones might include: • Initial project meeting to define the deployment scope • Purchase of EMC Smarts software • Project development begins • Installation in test environment complete • Testing complete • Installation in production environment complete • EMC Smarts goes live Additional information on scheduling is beyond the scope of this guide.

Architectural information checklist

Use the following checklist to aid in gathering information for your architectural design.

Table C-2 Architectural information checklist

Complete	Task	Descripiton	Related documentation
❑	Describe the organization's requirements and expectations.	Organization's vertical market: _____ (Reference to an organization's documentation) _____ _____ _____	"Determine the organization's requirements" on page 2-2
❑	Obtain network diagrams.	Ensure the diagrams include the locations of the following: • Network Operations Center (NOC) and LANs • Routing and switching devices • Firewalls • WAN links • High speed network technologies such as FDDI and Fast or Gigabit Ethernet In addition, important IP addresses and address ranges should be indicated.	"Obtain network information" on page 2-3
❑	If possible, schedule and discover the network.	If you intend to use the IP Discovery Manager to inventory the organization's network, schedule a time to perform the process and then discover the network as scheduled.	"Obtain network information" on page 2-3
❑	Describe the organization's network priorities.	Document these priorities in the Deployment Build Guide.	"Network priorities" on page 2-4
❑	Get the organization's testing/acceptance requirements.	Your design may be required to meet test and acceptance requirements. Obtain any specifications that cover integration testing, user acceptance testing, and operational acceptance testing. You may be required to write an installation or deployment report that follows an organization's particular standards.	"Determine requirements for installing software" on page 2-5
❑	Describe the organization's requirements for installing new software.	❑ Lab installation and testing ❑ Staging (*strongly* recommended) ❑ Preproduction deployment ❑ Shadow operation period (existing MoM still used) ❑ Other _____ Document these requirements and how the design meets them in the Deployment Build Guide.	"Determine requirements for installing software" on page 2-5
❑	List the products that currently monitor the network and will be integrated with EMC Smarts.	EMC Smarts' open architecture allows easy integration with third-party software. Many networks have at least a rudimentary network availability monitoring or a legacy system such as HP OpenView or IBM NetView. Document the products (including version) in deployment Build Guide.	"Integrating existing software with EMC Smarts software" on page 2-5
❑	List device types to manage.	To ensure devices are certified in EMC Smarts, obtain a list of the manufacturers and models for all devices in the network. Document the types of managed devices in the Deployment Build Guide.	"Identify the types of equipment in the network" on page 2-5
❑	Determine the number of managed ports and interfaces in the network.	Document all quantities and calculations used to determine the number of managed ports and interfaces in the Deployment Build Guide.	"Determine number of managed network devices" on page 2-6
❑	Estimate potential growth in quantity of managed devices.	The EMC Smarts deployment must support potential network growth. Estimate the growth over a specific time period. Document the calculations in the Deployment Build Guide.	"Accounting for network growth" on page 2-9

Table C-2 Architectural information checklist (continued)

Complete	Task	Descripiton	Related documentation
❑	Estimate number of managed systems and network adapters for licensing.	The EMC Smarts deployment can only discover and manage the quantity of systems and network adapters that are licensed. Document the quantities in the Deployment Build Guide.	"Determine quantities of devices for licensing" on page 2-10
❑	Describe the network security.	Describe security features such as the firewalls that will be between parts of the EMC Smarts deployment and if access lists are used. Obtain SNMP security parameter values for each device where they are used: for SNMP V1 and V2C, obtain community strings; for SNMP V3, obtain the user name, SNMP engine ID (optional), authentication protocol and password, privacy protocol and password (currently EMC Corporation supports AES and DES privacy protocols), and context name, if used. Document the security features in the Deployment Build Guide.	"Gather network security information" on page 2-11
❑	List any other network requirements or features that may affect EMC Smarts.	Document the features in the Deployment Build Guide.	"What other network features affect EMC Smarts?" on page 2-12

Solution architecture diagram checklist

Use the Solution Architecture Diagram to document your initial overall design of the EMC Smarts deployment.

Table C-3 Solution architecture diagram checklist

Complete	Task	Description	Related documentation
❑	List important device quantities on the Solution Architecture Diagram and in the Deployment Build Guide.	Start the Solution Architecture Diagram by listing the totals for Routers, Switches, Hubs, Bridges, Hosts, Ports and Interfaces. Include expected growth rate and estimates for managed ports and interfaces. Also document the quantities in the Deployment Build Guide.	"Documenting the deployment" on page 3-2
❑	Choose a tier size for platforms supporting Domain Managers.	❑Small (P&I:10K/AM or 5K/AM PM) ❑Medium (P&I:25K/AM or 15K/AM PM) ❑Large:(P&I:50K/AM or 25K/AM PM) ❑Very Large:(P&I:100K/AM or 50K /AM PM) Document the choice in the Deployment Build Guide.	"Determine the required size of the EMC Smarts deployment" on page 3-3
❑	Determine the quantity of Domain Managers required.	Include representations on the Solution Architecture Diagram.	"Determine the required size of the EMC Smarts deployment" on page 3-3
❑	Document equipment supporting EMC Smarts components.	Document the hardware and operating system supporting each of the EMC Smarts components. Start a chart for each host in the Deployment Build Guide and list the EMC Smarts components on the host.	"Determine the required size of the EMC Smarts deployment" on page 3-3
❑	Locate the hosts supporting the EMC Smarts components.	Document choices on the Solution Architecture Diagram and in the Deployment Build Guide.	"Locating EMC Smarts managers and platforms" on page 3-7
❑	Determine license server and licensing configuration requirements.	Document requirements on the Solution Architecture Diagram and in the Deployment Build Guide.	"Considering volume licensing configurations" on page 3-9
❑	Determine security requirements.	Document requirements in the Deployment Build Guide.	"Considering security and firewalls" on page 3-9
❑	Is Failover Capability Required for EMC Smarts?	❑No ❑Yes: Contact EMC Professional Services. Document choices on the Solution Architecture Diagram and in the Deployment Build Guide.	"Considering high availability configurations" on page 3-10
❑	Determine if overlapping IP networks are used.	Document needs in the Deployment Build Guide.	"Designing for overlapping (duplicate) IP networks" on page 3-11
❑	Plan acceptance tests and completion criteria.	Document in the Deployment Build Guide as each portion of EMC Smarts functionality is designed. Use them in validation.	"Designing acceptance tests" on page 3-12

Discovery design checklist

Before discovering the network using EMC Smarts, the requirements in the following checklist must be completed.

Table C-4 Discovery design checklist

Complete	Task	Description	Related documentation
Initial Discovery			
❑	Define a method for the initial topology discovery.	❑ Use a comprehensive seed file without autodiscovery. ❑ Use autodiscovery with an agent or seed file. Document the method in the Deployment Build Guide.	"Initial discovery" on page 4-2
Topology Maintenance and Subsequent Discovery			
❑	Define a schedule for Full Discovery.	Define a regular schedule for full discovery. Choose a time of relative inactivity. Document the schedule in the Deployment Build Guide. Include *crontab* or *sm_sched* control file entries if used.	"Topology maintenance and subsequent discovery" on page 4-4
❑	Define a schedule for Pending Discovery.	Define a regular schedule for pending discovery. Choose a time of relative inactivity. Document the schedule in the Deployment Build Guide. Include *crontab* or *sm_sched* control file entries if these utilities are used.	"Topology maintenance and subsequent discovery" on page 4-4
❑	Determine if Autodiscovery is appropriate.	Document choice in the Deployment Build Guide.	"Topology maintenance and subsequent discovery" on page 4-4
❑	Choose a method for adding devices to the topology.	❑ Agent without autodiscovery. ❑ Seed file without autodiscovery. ❑ Use autodiscovery with an agent or seed file. Document choice in the Deployment Build Guide.	"Adding new systems to an existing topology" on page 4-5
❑	Prepare Seed File or Choose Agent.	If a seed file will be used to add devices to the topology, obtain a list of devices with names or IP addresses. Document how to obtain the list or the location of the list in the Deployment Build Guide. If an agent will be used instead, document the IP address or name of the agent.	"Adding new systems to an existing topology" on page 4-5
❑	Define Autodiscovery Filters.	If autodiscovery is enabled, configure autodiscovery filters. These are inclusive filters that add devices to the topology. Document the autodiscovery filter criteria in the Deployment Build Guide.	"Controlling autodiscovery with filters" on page 4-5
❑	Define an Exclude Filter.	To exclude specific devices, use the exclude filter in the *discovery.conf* file. This simplifies creation of the autodiscovery filters. Document exclude filter entries in the Deployment Build Guide.	"Controlling autodiscovery with filters" on page 4-5
❑	Obtain SNMP security parameters per device.	Domain Managers use SNMP to poll the device agents. In order to do this, the Domain Manager needs the appropriate security information for the SNMP version: V1 and V2C use READ community strings for every device that will be managed; V3 uses the user name, SNMP engine ID, authentication protocol and password, privacy protocol and password and the context name. These parameters will be needed during discovery. Document in the Deployment Build Guide if permitted.	"Discovery and security" on page 4-6

Table C-4 Discovery design checklist (continued)

Complete	Task	Description	Related documentation
❑	Open necessary firewall ports.	If there is a firewall between any portions of the management infrastructure, certain TCP and UDP ports in the firewall must be opened for proper communications: • SNMP polls: 161 • SNMP traps: 162 • Broker: 426 • License Manager: 1744 • Domain Managers (1 per manager): configurable • EMC Smarts adapters, including SNMP Trap Adapter (Receiver) and Syslog Adapter: configurable Document the ports that are opened in the Deployment Build Guide.	"Discovery and security" on page 4-6
❑	Provide access to network devices to manage.	For each device that EMC Smarts will monitor, the device's access list must include the IP address of the hosts where Domain Managers are installed. EMC Smarts must have full access to browse the MIBs of the devices. Document in the Deployment Build Guide.	"Discovery and security" on page 4-6
❑	Ensure DNS is properly configured.	For EMC Smarts to name devices in its topology correctly, the DNS needs to be clean (proper forward and reverse lookup). If DNS is not used, use of an */etc/hosts* file or not doing any name resolution at all can be considered.	"Discovery and the Domain Name System" on page 4-7
❑	Determine if Discovery post processing is required.	Determine if discovery post processing using ASL rulesets will be used. Document in the Deployment Build Guide.	"Advanced discovery post-processing" on page 4-7
❑	List unreachable IP Addresses	If there are groups of IP addresses that are NOT normally reachable, assemble a list of IP ranges or some matching criteria so EMC Smarts will not unnecessarily ping these addresses. Document these addresses in the Deployment Build Guide.	"Advanced discovery post-processing" on page 4-7

Polling and threshold checklist

Table C-5 Polling and threshold checklist

Complete	Task	Description	Related documentation
❑	Determine polling group requirements.	Design polling groups based on importance of network device performance both to the network and to the various parts of the organization. Also consider network latency to determine if changes are needed. Document choices in the Deployment Build Guide.	"Polling and polling groups" on page 5-2
❑	Set polling parameters for each polling group.	Set polling parameters based on importance of network device performance. Additional modifications may be necessary if polling does not present an accurate picture of network availability during validation. Document new polling parameters in the Deployment Build Guide.	"Polling and polling groups" on page 5-2
❑	Determine threshold group requirements.	Design threshold groups based on importance of network device performance both to the network and to the various parts of the organization. Document choices in the Deployment Build Guide.	"Threshold values and threshold groups" on page 5-3
❑	Set threshold parameters for each threshold group.	Set threshold parameters based on the expected effect of degraded performance on network operations. Additional modifications may be necessary during validation and as the organization gains experience with the performance indicators. Document new threshold parameters in the Deployment Build Guide.	"Threshold values and threshold groups" on page 5-3

Syslog processing checklist (optional)

Table C-6 Syslog processing checklist

Complete	Task	Description	Related documentation
❑	Determine how to create the file for processing by the Syslog Adapter.	A file must be created that the Syslog Adapter can parse. Determine which devices will contribute messages to the file. Consistent layout of the messages in the file is required for Syslog Adapter processing. Include all details in the Deployment Build Guide.	"Creating the syslog file" on page 6-2
❑	Determine the location of the file that the Syslog Adapter will process.	The process that is creating the file must be able to receive messages from source applications and the created file must be accessible by the Syslog Adapter. Include all details in the Deployment Build Guide, including location, host and path.	"Processing the syslog file" on page 6-2
❑	Choose the messages that are most important for processing.	Choose the messages that are most important for processing. Include all details in the Deployment Build Guide.	"Processing the syslog file" on page 6-2
❑	Determine the characteristics of the notifications that are generated.	For each message that generates a notification, determine the notification format. These characteristics will be used to develop the hook script for the EMC Smarts syslog processing deployment. Include all details in the Deployment Build Guide.	"Processing the syslog file" on page 6-2
❑	Add Syslog Processing to the Solution Architecture Diagram.	Add Syslog Processing to the Solution Architecture Diagram.	"Processing the syslog file" on page 6-2

Index

A

Access lists 2-11, 4-6
Adapter
    Syslog 6-1
Adapters 2-5
Administration
    Design users 7-1
Administration Design Checklist 7-9
Architectural Information Checklist 2-13
Architecture 1-2
Autodiscovery 4-3
    Filtering 4-3, 4-5
    Initial discovery 4-3

B

Backup Interface Support Thresholds 5-3
BASEDIR viii
Batching Traps
    Performance improvements 10-11
Before You Begin Checklist 1-7, C-2
brcontrol 9-3

C

Certifications 2-5
Certified Devices 4-6
Checklists C-1
    Administration Design 7-9
    Architectural Information 2-13
    Before You Begin 1-7, C-2
    Discovery Design 4-8, C-6
    Polling and Threshold 5-4
    Solution Architecture Diagram 3-13, C-5
Cisco Discovery Protocol (CDP) 4-3
clientConnect.conf 7-3
Codebook Tasks 10-5
Community Strings 2-11, 4-6
Consoles
    Defining 7-6, 7-7
    Restricting operations 7-6
cron 4-4
Customer support xi

D

Deployment
    Design description 1-4
    Installation and configuration description 1-5
    Phases 1-4
    Process 1-4
    Tuning description 1-6
    Validation description 1-6
Devices
    Certifications 2-5
    Estimating network growth 2-9
    Estimating quantity of network 2-6
Dial-on-Demand Interface Support Thresholds 5-3
Discovery
    Advanced post-Processing 4-7
    Autodiscovery 4-3
    Autodiscovery Filtering 4-3
    Definition of 4-2
    Duration 10-6
    ICMP B-3
    Initial 4-2
    Maintenance of topology 4-4
    Performance 10-6
    Post Processing Performance 10-6
    Scheduling 4-4
    Security 4-6
    Seed file 4-2
    SNMP B-3
        Traps B-3, B-6
    Validation 9-4
Discovery Design Checklists 4-8, C-6
discovery.conf 4-3, 4-5
dmctl A-5
Domain Name System (DNS) and Discovery 4-7

E

EMC customer service xi
EMC product information xi
EMC technical support xi
Extreme Discovery Protocol (EDP) 4-3

F
Failover capabilities 2-12
Filters
    Autodiscovery 4-5
Firewalls 2-11
    Ports and Discovery 4-6
Foundry Discovery Protocol (FDP) 4-3
Full Discovery 4-4

G
Global Console
    Performance 10-11
Global Console Users
    Categories 7-2
    Security 7-2
Growth, Estimating Network 2-9

H
Help xi

I
ICMP B-3
ICMP Statistics 10-5, 10-8
Initial Discovery 4-2
    Autodiscovery 4-3
Integrating Existing Software 2-5
IP addresses
    Overlapping 3-11, B-1
IP Management Domain B-2
ipExcludeList 4-5

L
Licensing 2-6, 2-10
    Configurations 3-9
    License Manager communications port 4-6
    Permanent 8-3

M
Memory
    Process size limit 10-3
my_hook_syslog.asl 6-2

N
name-param.conf 4-7
Network Devices
    Certified 4-6
    Estimating growth 2-9
    Estimating quantity 2-6
    Types 2-5
Network Diagrams 2-3
Network Security 2-11
    Firewalls 2-11
Notifications
    Lists 7-5

O
Operating System Resource Limits 10-3
Out-of-Band Management 2-12
Overlapping IP addresses 3-11, B-1
    IP Management Domain B-2

P
Password Configuration 7-3
Pending Discovery 4-4
Performance
    Discovery 10-6
    Global Console 10-11
    Improving 10-10
Performance Tuning
    Guidelines 10-2
Policy-based routing B-4
Polling 5-1
    Groups 5-2
    Validation 9-5
Polling and Threshold Checklist 5-4
Post Processing Discovery Performance 10-6
Process Size Memory Limit 10-3
ps 10-3

R
Reconfiguration 10-7
Repository, Saving 10-7
Requests for Proposal (RFPs) 2-2, 2-3, 2-6
Resource Requirements 10-3
Router
    Policy-based routing B-4
    Source routing B-5

S
Saving the Repository 10-7
Scripts and naming conventions 2-3
Secret Phrase 8-5
Security 2-11
    clientConnect.conf 7-3
    Community strings 4-6
    Discovery 4-6
    Global Console Users 7-2
    serverConnect.conf 7-3
    SNMP V3 4-6
Seed File 4-2
serverConnect.conf 7-3
Service xi
sm_sched 4-4
sm_snmp 9-6
sm_snmpwalk 2-5
sm_topodump A-5
sm_tpmgr 4-4, 10-4, 10-5, 10-6, A-5
SmartPacks 2-5
SNMP B-3
    Community strings 2-11, 4-6
    SNMP V3 and Autodiscovery 4-3
    SNMP V3 security 4-6
    Statistics 10-5

Traps B-3, B-6
SNMP Network Security
SNMP community strings 2-11, 4-6
SNMP Statistics 10-8
SNMP V3
See SNMP
Software
Integrating Third-party 2-5
Solution Architecture Diagram Checklist 3-13, C-5
Source Routing B-5
Split Topology 10-10
Syslog (System Log) File 6-1
Syslog Adapter
Processing 6-2
Tailing 6-2
Template 6-2

T

Tailing, Syslog 6-2
Technical support xi
Third-Party Software
Discovery and 4-2
Threshold Groups 5-3
Thresholds 5-1
Adjusting 10-11
TM_USESEEDNAME 4-7
Tool Deployment 7-8
Trap Processing
Batching 10-11
Configuration 8-6
Deployment procedure 8-6
SNMP B-6
Validation 9-6
trap_mgr.conf 9-6

U

User Profiles 7-4, 9-7
Validation 9-7
Users
Deploy Configurations 8-7
Validation 9-7

V

Validation 9-2
Vertical Markets 2-2
Virtual Routers 2-7
Licensing 2-10