Configuring SGM to Run with Various Networking Options

In addition to running on standard IP-connected networks, SGM has the flexibility to adapt to a variety of different networking environments, including Virtual Private Network (VPN), Network Address Translation (NAT), firewall, port-forwarding, and Secure Sockets Layer (SSL). SGM can run in each of these environments individually, or in any combination of networking environments.

This appendix describes communication between the SGM client and the SGM server. As shown in Figure D-1, this includes:

- Two-way RMI communication between a Java-based GUI client and Java-based server processes. The client can send requests to and receive responses from the server, and the server can send unsolicited notifications to the client. For example, if the server detects that a router's state has changed, it sends a notification to all SGM clients to update their Topology windows.
- One-way HTTP communication between a Web browser and an SGM-embedded Web server, using the request/response model.

Figure D-1 SGM Communication

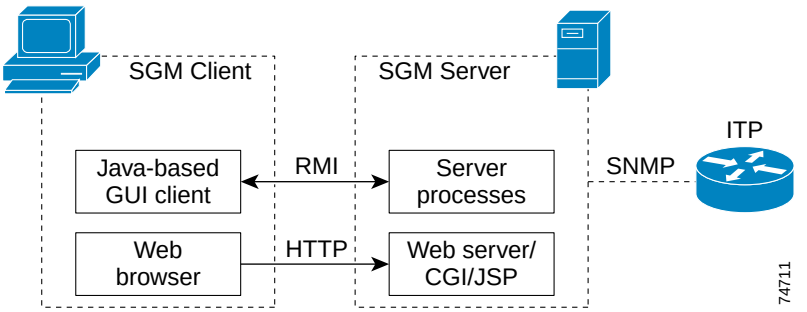

Note This appendix does not address communication between the SGM server and the router, which uses the SNMP protocol for network management.

This appendix includes the following sections:

VPN Communication

SGM client/server communication can run transparently through a VPN tunnel, which is a secure IP layer, without any user intervention. You can use VPN to connect to a corporate network, then start the SGM client to connect through the VPN tunnel to an SGM server in the corporate network.

When the client host establishes a VPN tunnel, the operating system (or system library) sees this as another virtual IP interface. The VPN tunnel does not affect HTTP communication between the Web browser and server, it only affects RMI communication between the SGM client and server processes.

For HTTP communication, the virtual IP address is transparent to the upper layer. The operating system automatically chooses the correct IP address to send out the request packet. For RMI communication, the SGM client needs to register with the SGM server server using the correct IP address, so that the server can invoke RMI callbacks and send unsolicited notifications to the client.

SGM solves this problem by automatically detecting the local IP interface so that the SGM server can send unsolicited notification to the correct IP address.

Figure D-2 shows a sample VPN network with the following characteristics:

- The SGM client with IP address 192.168.0.1 is connected to the SGM server network through a VPN tunnel.
- The SGM client host has obtained VPN IP address 10.1.1.2, which is a virtual IP interface.

Figure D-2 VPN Communication

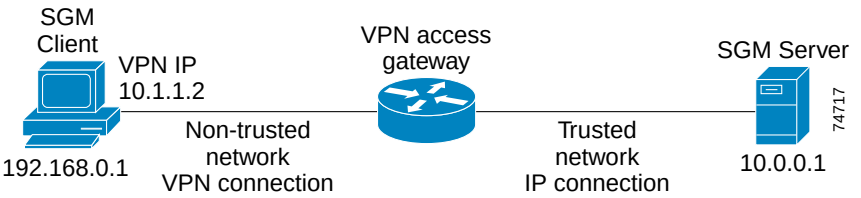

When connecting to the SGM server, the SGM client automatically recognizes its VPN IP address, 10.1.1.2, and uses that address to register with the SGM server to receive RMI callbacks. This configuration is transparent to the user; no manual configuration is needed.

NAT Communication

SGM client/server communication can run through one or more static NAT-connected networks. (SGM does not support dynamic NAT or dynamic NAT pool overloading.)

In a static NAT network, the SGM client and server are located on different sides of the NAT network, with no routes between the client network and the server network. The NAT device statically maps the client IP address to a NAT address in the server network, and the server IP address to a NAT address in the client network.

The NAT device translates packets between the SGM client and server by replacing IP address headers when packets pass through. From the client's point of view, the server appears to be at a NAT IP address in the client network, and vice versa. For most protocols, this technique is sufficient to enable the client and server to communicate.

However, for RMI protocol, this is not sufficient. RMI protocol requires the client and server to keep remote object references by remote stubs. These remote stubs contain the remote objects' IP addresses, and are passed between the client and server using Java serialization. The NAT device only converts the IP addresses in the IP packet header, but the remote stub object is within the packet content. Therefore, the NAT device cannot recognize the IP address inside the packet, and fails to route it correctly.

SGM solves this problem by creating a specialized NAT-aware socket factory. Some manual configuration on the part of the user is required to enable SGM to “know” the network NAT configuration.

Figure D-3 shows a sample static NAT network with the following characteristics:

- A static NAT device connects Network A (192.168.*.*) to Network B (10.*.*.*), with no routes between Network A and Network B.
- The NAT device maps SGM client IP address 192.168.0.1 in Network A to 10.1.1.2 in Network B.
- The NAT device maps SGM server IP address 10.0.0.1 in Network B to 192.168.1.2 in Network A.

Figure D-3 Static NAT Communication

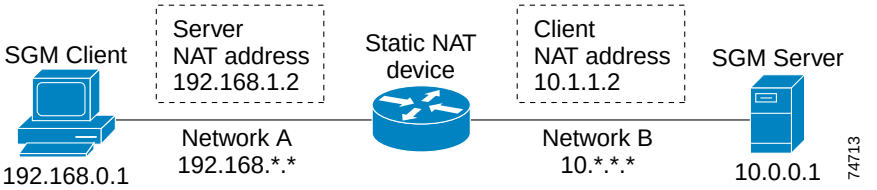

To configure SGM in this static NAT network, you must modify the SGM client's *RMIOverNAT.properties* file.

- If you installed SGM in the default directory, */opt*, then the location of the file is */opt/properties/RMIOverNAT.properties*.
- If you installed SGM in a different directory, then the file is located in that directory.

For the example shown in Figure D-3, you must add the following line to the file:

10.0.0.1 = 192.168.1.2

This line maps the SGM server's real IP address, 10.0.0.1 in Network B, to its NAT address, 192.168.1.2, in Network A, which is the server's IP address as seen by the client.

Note The SGM server automatically detects the SGM client's NAT address. No manual configuration on the part of the user is needed at the server side.

Firewall Communication

To enable SGM client/server communication through a firewall, you need to set up the firewall so that it allows SGM communication packets to pass through freely.

Note The SGM client and server communicate using TCP sockets. All port numbers in this section are TCP ports.

The port number used by SGM is configured in the *System.properties* file:

- If you installed SGM in the default directory, */opt*, then the location of the file is */opt/properties/System.properties*.
- If you installed SGM in a different directory, then the file is located in that directory.

Set the following parameters on the server side of the file:

RMIREGISTRY_PORT = 44742

DATASERVER_PORT = 0

MLSERVER_PORT = 0

PMSERVER_PORT = 0

WEB_PORT = 1744

where:

- **RMIREGISTRY_PORT** is the port on which the RMI naming server listens. You must specify a port number; **0** is not allowed.
- **DATASERVER_PORT** is the port on which the sgmDataServer process listens. If you specify **0**, SGM uses any available port, 1024 and above.
- **MLSERVER_PORT** is the port on which the sgmMsgLogServer process listens. If you specify **0**, SGM uses any available port, 1024 and above.
- **PMSERVER_PORT** is the port on which the sgmProcMgrServer process listens. If you specify **0**, SGM uses any available port, 1024 and above.
- **WEB_PORT** is the port on which the SGM Web server listens. You must specify a port number; **0** is not allowed. To change the **WEB_PORT** number, use the **sgm webport** command. See the "SGM Commands and Descriptions" section on page B-2 for more information on the use of this command.

If any of these port numbers changes, you must restart the SGM server before the changes take effect.

Set the following parameters on the client side of the file:

RMIREGISTRY_PORT = 44742

CLIENT_PORT = 0

where:

- **RMIREGISTRY_PORT** is the port on which the server-side RMI naming server listens. This port number must match the one specified for the **RMIREGISTRY_PORT** on the server side.
- **CLIENT_PORT** is the port on which the SGM client listens for RMI callbacks (unsolicited notifications). If you specify **0**, SGM uses any available port, 1024 and above. If you specify **CLIENT_PORT** with a value other than **0**, you can run only one SGM client process at a time.

If any of these port numbers changes, you must restart the SGM client before the changes take effect.

Figure D-4 shows a sample firewall network with the following parameters set in the *System.properties* file:

- On the SGM server side:

  RMIREGISTRY_PORT = 44742

  DATASERVER_PORT = 44751

  MLSERVER_PORT = 44752

  PMSERVER_PORT = 44753

  WEB_PORT = 1774

- On the SGM client side:

  RMIREGISTRY_PORT = 44742

  CLIENT_PORT = 56173

Figure D-4 Firewall Communication

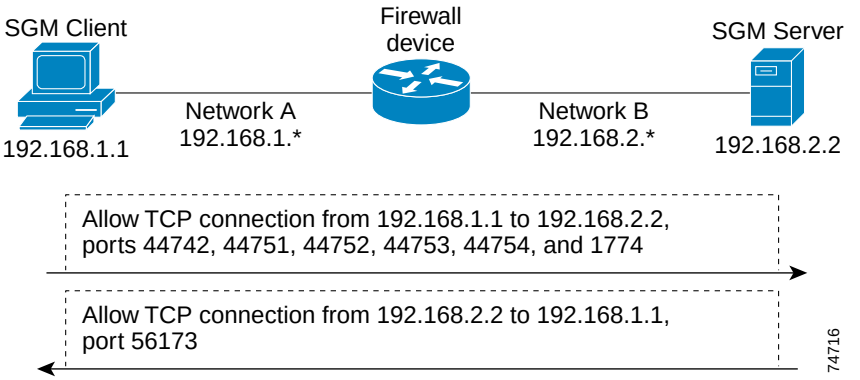

Port-Forwarding Communication

To enable SGM to operate in a TCP port-forwarding environment, perform the following configuration tasks:

Step 1 Configure the server hostname and port number mapping in the SGM client's *RMIOverNAT.properties* file, as described in the "NAT Communication" section on page D-4.

Step 2 Configure the port numbers used by the SGM client and server in the *System.properties* file, as described in the "Firewall Communication" section on page D-5.

Step 3 Configure the port-forwarding tunnel to forward each side's TCP connection to the other side.

Figure D-5 shows a sample network that uses Secure Shell (SSH) port-forwarding. Other port-forwarding configurations might use a single host with dual interfaces at the client's and server's networks. While other port-forwarding configurations may differ from this example, the general rules to configure SGM to operate in a port-forwarding environment are the same.

Figure D-5 Port-Forwarding Communication

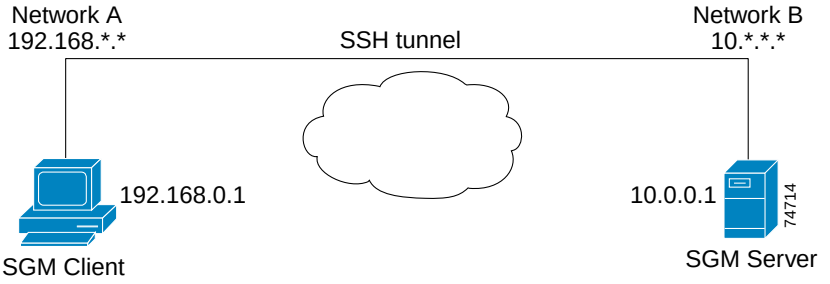

The port-forwarding network shown in Figure D-5 has the following parameters set;

- In the *System.properties* file, on the SGM server side:

  RMIREGISTRY_PORT = 44742

  DATASERVER_PORT = 44751

  MLSERVER_PORT = 44752

  PMSERVER_PORT = 44753

  WEB_PORT = 1774

- In the *System.properties* file, on the SGM client side:

  RMIREGISTRY_PORT = 44742

  CLIENT_PORT = 56173

- In the SGM client's *RMIOverNAT.properties* file:

  10.0.0.1/44742 = 192.168.1.2/25742

  10.0.0.1/44751 = 192.168.1.2/25751

  10.0.0.1/44752 = 192.168.1.2/25752

  10.0.0.1/44753 = 192.168.1.2/25753

  10.0.0.1/1774 = 192.168.1.2/8080

- In the port-forwarding network:

  Local port 25751 => remote host 10.0.0.1, port 44742

  Local port 25751 => remote host 10.0.0.1, port 44751

  Local port 25752 => remote host 10.0.0.1, port 44752

  Local port 25753 => remote host 10.0.0.1, port 44753

  Local port 8080 => remote host 10.0.0.1, port 1774

  Remote port 56173 => local host 192.168.0.1, port 56173

Note For port-forwarding setup, the backward-forwarding port numbers must match each other. In the above example, both are **56173**. The forward-forwarding port number do not need to match each other.

SSL Communication

If SSL is implemented and enabled in your SGM system, SGM uses secure socket communication for both RMI and HTTP communication between the SGM client and server.

SGM supports standard-based SSL encryption algorithms, including RSA, DSA public key algorithms, and 40-bit or 128-bit encryption. SGM can generate an X.509 certificate and a certificate signing request (CSR), which is interoperable with most certificate authorities (CAs).

Both the SGM Web server and the SGM server processes share the same SSL key/certificate pair. Both the SGM client and the Web browser can examine the server's certificate.

For more information, including descriptions of the SGM commands and procedures used to implement, enable, manage, and monitor SSL support, see the "Implementing SSL Support in SGM" section on page 4-24.

Figure D-6 shows a sample SGM-over-SSL network with the following characteristics:

- A user-generated SSL key pair on the SGM server.
- The server's certificate is trusted on the SGM client.
- Communication between the client and server is RMI-over-SSL and HTTPS. Both protocols are encrypted and secure.

Figure D-6 SSL Communication