Common Firewall Traversal Issues

Previous

Next

Let's assume that a client and server need to communicate over an untrusted network, and that the client and server hosts reside in private networks behind firewalls:

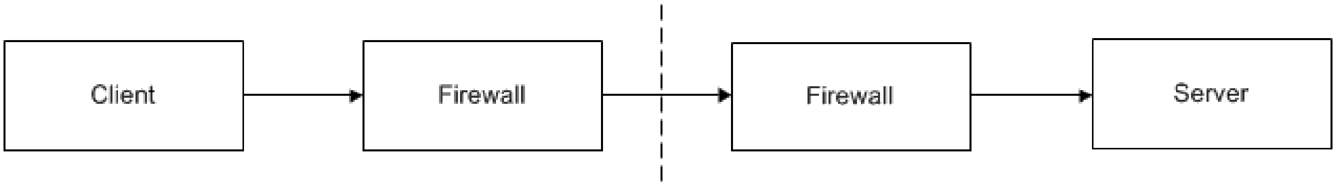

Scenario 1: Client request in a typical network.

Although the diagram looks fairly straightforward, there are several troublesome issues:

- A dedicated port on the server's firewall must be opened and configured to forward messages to the server.
- If the server uses multiple endpoints (e.g., to support both TCP and SSL), then a firewall port must be dedicated to each endpoint.
- The client's proxy must be configured to use the server's "public" endpoint, which is the host name and dedicated port of the firewall.
- If the server returns a proxy as the result of a request, the proxy must not contain the server's private endpoint because that endpoint is inaccessible to the client.

To complicate the scenario even further, the illustration below adds a callback from the server to the client. Callbacks imply that the client is also a server, therefore all of the issues associated with previous illustration now apply to the client as well.

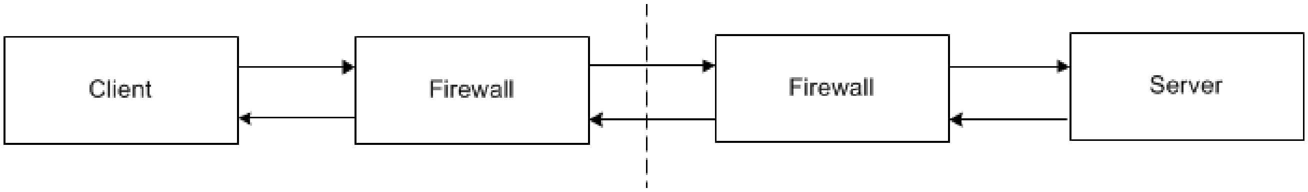

Scenario 2: Callbacks in a typical network.

As if this was not complicated enough already, the illustration below adds multiple clients and servers. Each additional server (including clients requiring callbacks) adds more work for the firewall administrator as more ports are dedicated to forwarding requests.

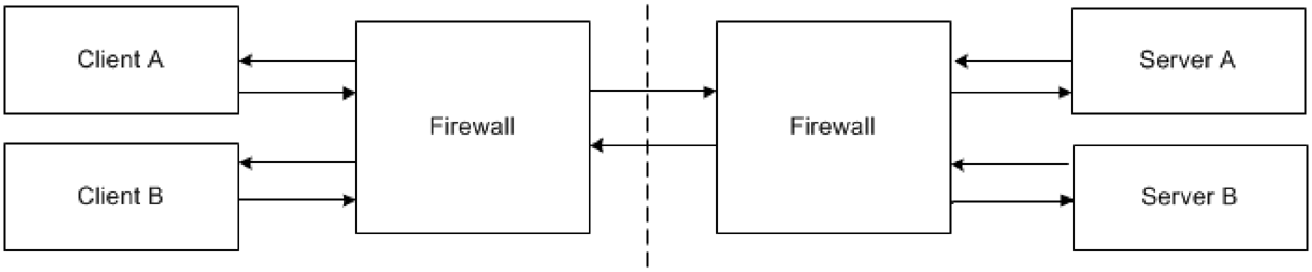

Scenario 3: Multiple clients and servers with callbacks in a typical network.

Clearly, these scenarios do not scale well, and are unnecessarily complex. Fortunately, Ice provides a solution in Glacier2.

Back to Top ^

See Also

- About Glacier2
- How Glacier2 Works
- Getting Started with Glacier2

Previous

Next