

JMeterRemoteTestingEnhancements

Navigation trail: [JMeterProjectPages](#) - [JMeterDevelopment](#) - [JMeterDevelopment/LongTerm](#)

How can we improve the remote testing experience?

- Documentation of remote testing does not provide details of the ports that need to be open on a firewall
- Make remote testing easier and less error prone (provide shell scripts for rmi registry or do not use RMI at all) – [OliverRossmueller](#) 2003-01-06
 - The very first thing we should do is to update the jmeter-server scripts to set the classpath and start rmi registry before starting jmeter. – [ScottEade](#) 2003-01-07
- Need a solution for tunneling RMI through a firewall (publicly available solutions exist). – [ScottEade](#) 2003-01-05
- Option to specify the port to be allocated by rmi. – [ScottEade](#) 2003-01-05 *I just can't understand any situation where running JMeter remotely through a firewall is necessary or desirable. I can only think that people who insist on being able to do this have a fundamental misunderstanding of what a remote JMeter "server" is for* – [MikeStover](#) _ Situation: I want to run the jmeter GUI from my desktop in our local network, but test an internal service in our data center. The remote jmeter slaves and our server need to be in the data center network to simulate real usage (The users of this service are internal). If I ran the tests from my local network there would be extra latency and it would not give me realistic numbers. – [MikeDeeks](#) _)
- The ability to easily configure a set of remote testing hosts with varying parameters (number of threads to cater for testing hosts of varying capacity, counter values to handle things like ensuring user identifiers are unique across all test hosts) so that it is easier to initiate distributed testing from multiple remote hosts. – [ScottEade](#) 2003-01-07
 - Ensure relevant items can be set using variables.
 - Allow parameters to be specified on a per server basis (much like User Parameters), perhaps these would replace the User Defined Variables on the Test Plan element.
 - Would want a checkbox to indicate which servers should be active for a test run (this would effectively define a single set of hosts, the members of which can be very easily updated).
 - Could enter the address of the server here and eliminate the remote_hosts setting in jmeter.properties.
 - Good thoughts - maybe a dialog box could provide a bunch of controls associated with starting a remote host - ie IP address, custom user-defined variables for that run, etc.
- Pass properties specified via jmeter -Jvariable=value through to remote servers. It took embarrassingly long for me to realize that if I pass -J properties on the command line and don't run with -r, the properties are set correctly for the local instance, but if I do pass -r, none of the remote servers see the property specified on the command line (they always get the default value). --ChrisNuzum 2006-10-18
 - I also took a long time to resolve this issue... On the client, code -Jmygroup.threads=10 and -Gmygroup.threads=10. -G parameters get passed to the server. --Dave Bruce 2009-12-10
- Provide a simple way to make a counter work with multiple machines. A lot of use cases the only data that needs to be changed is the usernames and deploying a csv file and all is unnecessary and inconvenient. I've actually written a component that would do this, but the number of servers being used and an index in the list are required to do this properly. Should be very easy to do this when initializing engines. just define two more standard test variables in each engine as created (Sean Ferguson email at sean.d.ferguson@att.net)

RMI/Firewall issues

In the list above are some points regarding firewall configuration and the need for tunneling RMI through a firewall. To make these requirements more clear can anyone please give an example of why it is necessary to use a remote JMeter server BEHIND a firewall and why it is not possible in that situation to use a JMeter server IN FRONT of the firewall. – [OliverRossmueller](#)

Because the firewall is placed between the admin network (where I'm sitting) and the production network (where the servers are). Of course there are other solutions, such as running JMeter GUI on the production network displaying on my screen through an X tunnel – but these are often undesirable (because of latency/bandwidth limitations) or impossible (because some machines don't support X). This is nothing weird nor unusual, I'd say I'd find myself in such situation in 70-80% of the large customers I visit – [JordiSalvat](#) 2003-08-11.

I noted the firewall issues when I was doing something a little weird - using a win2k box to control a JMeter server on a Mac OS X box with it's firewall running (even though it was on the internal network). The Mac is a laptop and so it is sometimes connected directly to the net, hence the local firewall. – [ScottEade](#) 2003-01-16

- It should be possible to deal with this setup by adding some firewall rules to allow RMI connections to the Mac from the local network, so from my point of view this is a configuration issue and not a JMeter issue. – [OliverRossmueller](#)
 - Yes, but it would help if JMeter provided documentation as to what ports need to be open, or allowed the port to be specified, or provided a tunneling mechanism 😊 – [ScottEade](#) 2003-01-16
 - It is possible to get RMI work with two specified ports. The port used by the rmi registry is 1099 or any other port specified when starting the registry, so you can control this. The problem with RMI in combination with a firewall is the port used by the RMI server which by default is a random port chosen by the operating system. To get RMI through a firewall you therefore have to open all ports above 1023 because you do not know in advance which port will be used. Fortunately it is possible also to specify this port by replacing the default socket factory used by RMI with your own implementation which always uses the same port. This way you have to open only two ports on the firewall the rmi registry port and the JMeter server port. – [OliverRossmueller](#)

A different example might be that you have a bunch of JMeter server machines on an internal network somewhere that you want to control from a physically remote machine that is on an external network - this is a contrived example as it would require effort to make the JMeter server machines visible outside the firewall. – [ScottEade](#) 2003-01-16

- This might be an interesting and usefull feature in some situations. One possible way to do this is via web interface, something like a browser-based admin console to control your JMeter server instances. – [OliverRossmueller](#)
 - So this provides the example you requested above? Why not just control them through an enhanced remote server manager UI within JMeter itself (basically what this page is all about) and tunnel the RMI. – [ScottEade](#) 2003-01-16

- I was not sure why you would need RMI tunneling, but this scenario makes sense to me and RMI tunneling is one way to make it work. Do you know of any RMI tunneling implementation that fits with the Apache Licenses so we could use it for JMeter? – [Oli verRossmueller](#)

All in all it is nothing like the issue I originally imagined it to be. – [ScottEade](#) 2003-01-16

a modified version

re remote testing and firewall issues, just modified jmeter source based on v2.1.1. with the modification, a new property 'server.service.port' could be used for the jmeter service port (pls note this is different from the rmi registry port which could also be configured by 'server.rmi.port') in jmeter.properties. current jmeter release uses a random port number which may cause some trouble for sysadmin to adjust firewall rules (especially if you sysadmin is very much concerned about any potential security issues). so you could give jmeter a static port number for the service and let sysadmin know the port as well.

- if you only need binary (working copy), download ApacheJMeter_core.jar http://wiki.apache.org/jakarta-jmeter-data/attachments/JMeterRemoteTestingEnhancements/attachments/ApacheJMeter_core.jar which is recompiled with my change; and as an option, feel free to download new jmeter.properties <http://wiki.apache.org/jakarta-jmeter-data/attachments/JMeterRemoteTestingEnhancements/attachments/jmeter.properties>.
- for source code, pls download two patches, one for src/core/org/apache/jmeter/engine/RemoteJMeterEngineImpl.java <http://wiki.apache.org/jakarta-jmeter-data/attachments/JMeterRemoteTestingEnhancements/attachments/RemoteJMeterEngineImpl.java.patch.txt> and the other for bin/jmeter.properties <http://wiki.apache.org/jakarta-jmeter-data/attachments/JMeterRemoteTestingEnhancements/attachments/jmeter.properties.patch.txt>. simply apply them and job done.

(I have been looking for any chance contribute that modification back to main development tree, so feel free to use it if it could be included in a future release)

– [ChenWang](#) 2006-01-12

The above changes have been applied to the code - see Bug 30379

Remote Testing Set Up using Rendezvous

It would be really nice to use some of the [ZeroConf](#) technology that is become more and more proliferate now to automatically discover the remote servers available to control on a network. I know there is a module called JRendezvous which is an implementation of Rendezvous (Apple's brand name for [ZeroConf](#)) that could be used. Imagine how nice it would be to run up your remote servers, then click a menu in your client GUI and automatically see the various remote servers that are available. Its one last thing to fight with when trying to get this set up. Also looks a little more polished. ...although, saying this, it may be more useful if the GUI could cut and paste nodes first (let's run before we can walk) – Craig Pugsley 12/3/03

Remote Testing Set Up

Would it be feasible to have the remote server "push" its IP address to the controlling client?

– [KeivnHurley](#) 01/27/2005

Serialisation

At present, when a script is sent to the remote server for execution, AFAIK the entire tree is serialised and reconstructed at the other end. This places a lot of burden on developers to make sure that objects are serialisable. Since there is already a convenient representation for a test (i.e. the JMX format) it seems to me that we should make use of it and ship that across instead. It would mean that the remote installation would need to have all the appropriate classes present, but that is probably true anyway. It should reduce the amount of data sent. It would probably get rid of some of the incompatibilities that occur. I read somewhere recently (probably Effective Java by Joshua Bloch) that hash values are not *guaranteed* to be the same for different JVMs, nor even between runs of a JVM, which might explain some problems users have reported.[sebb](#)

Similar considerations apply to the Sample data that is returned, and it would be useful to be able to batch this and/or just send back aggregate data during a test run, and allow the client to request the full data when required.