



Consiglio Nazionale delle Ricerche

**EURid Benchmarking WP: Technical Report n.2.
Functional and Workload Stress Testing**

M. Bertoli, M. Buzzi, M. Conti, E. Gregori, C. Lucchesi,
M. Martinelli, G. Valente

IIT TR-26/2005

Technical Report

Dicembre 2005



Istituto di Informatica e Telematica

Abstract

This is a report of the workload stress and functionality tests performed on the EURid registration system.

The first part of our work was dedicated to design and develop an EPP testing software, and then we focused on the system functioning and performance. Our objective was to realize a suite of tests as scalable and reusable as possible.

As a preliminary testing we carried out functional server tests (as required to each EURid partners) discovering some minor fixes needed by the system, but our overall opinion is that it works fine from the functional point of view.

Concerning system performance, we carried out two sets of tests: a preliminary workload test was performed in September/October 2005 interacting with the “old” EPP system, the second set of workload stress testing was executed in November 2005 on the operational EURid EPP system.

Functioning Tests

Commands are idempotent (checked with login). OK

Wrong value of message length in an EPP message is not accepted. OK

Connection gets closed after a successful apply domain. OK

Connection gets closed after a successful create domain. OK

Contact Tests

Create contact

- it is possible to create several contacts using the same identical EPP command. Is it supposed to do so?
- maximum number of possible tech contacts for registrant (10) (EURid policy) (for the 11 tech). 2308 OK
- id size is 3-16 OK
- name size is 1-255 OK
- **org size declared 1-255 but is 0-255**
- street size is 1-255 OK
- city size is 1-255 OK
- **sp size declared 1-255 but is 0-255**
- **pc size declared 0-10 but is 1-16**
- only one billing contact per registrar allowed: return code (for the second billing): 2003 OK
- creation of a tech contact. OK
- creation of a registrant contact. There are no limits. 1000 OK
- creation of an Onsite contact. There are no limits. 1000 OK
- creation of a contact with an empty type. 2001 OK
- creation of a contact with a deprecated type licensee. 2001 OK
- creation of a contact with a deprecated type accmgr. 2001 OK
- **vat size declared 1-20 but is 0-20**
- languages permitted. 1000 OK
- create with lang empty. 1000 OK
- create with a language not permitted. 2001 OK

Delete contact

- id size 3-16 OK
- delete contact with an empty id. 2001 OK
- delete a nonexistent contact. 2303 Object does not exist OK
- delete a billing contact that has apply domain associated. 2305 Object association prohibits operation OK
- delete a billing contact that hasn't any apply domain associated. 1000 OK
- delete a tech contact that has apply domain associated. 2305 Object association prohibits operation OK
- delete a new tech contact that hasn't any apply domain associated. 1000 OK
- delete a registrant contact that has apply domain associated. 2305 Object association prohibits operation OK
- delete a registrant contact that hasn't any apply domain associated. 1000 OK
- delete a onsite contact that has apply domain associated. 2305 Object association prohibits operation OK
- delete a onsite contact that hasn't any apply domain associated. 1000 OK

Info contact

- An info contact with an empty id. 2001 OK
- An info contact with an empty clTRID. 2001 OK
- Retrieve information about a nonexistent contact. 2303 Object does not exist OK
- Retrieve information about the billing contact. 1000 OK
- Retrieve information about a tech contact. 1000 OK
- Retrieve information about a registrant contact. 1000 OK
- Retrieve information about a new registrant contact that hasn't any apply domain associated. 1000 OK
- Retrieve information about an onsite contact that has apply domain associated. 1000 OK
- Retrieve information about a new onsite contact that hasn't any apply domain associated. 1000 OK

Update contact

- Update an inexistent contact. 2303 OK
- **name size <contact:name> declared 1-255 but is 1-50**
- **org size <contact:org> declarated 1-255 but is 0-100**
- street size <contact:street> is 1-255 OK
- city size <contact:city> is 1-255 OK
- **sp size <contact:sp> declared 1-255 but is 0-255**
- **test pc size <contact:pc> declared 0-10 but is 1-16**
- update a generic field (contact:name) of a registrant contact that have applyDomain associated. 1000 OK
- create a new registrant (without domains associated) and update this registrant changing a generic field e.g: org 1000 OK
- update a generic field of a billing contact that have applyDomain associated. 1000 OK
- create a new onsite contact and update a generic field. 1000 OK
- update a generic field of a tech contact that have applyDomain associated. 1000 OK

- **test vat size declared 1-20 but is 0-20**
- test the languages permitted. 1000 OK
- test update with lang empty. 1000 OK
- test update with a language not permitted. 2001 OK

Domain Tests

Create domain

- **name size declared 1-255, but is 2-63**
- creation of a domain with an empty domain name. 2001 OK
- creation of a domain name with some capital letters. 1500 (the name is registered with minus letters) OK
- creation of a domain name with space digits. 2005 OK
- creation of a domain name with special chars. 2005 OK
- creation of a domain with reserved domain name. 2005 OK
- creation of a domain with dotted domain name. 2005 OK
- creation of a domain with double dotted domain name. 2005 OK
- creation of a domain with a TLD different from .eu. 2005 OK
- creation of a domain with a second level domain and .eu. 2005 OK
- creation of a domain with .eu. 1500 OK
- creation of a domain without .eu. 1500 OK
- creation of a domain with empty registrant contact. 2001 OK
- creation of a domain with a nonexistent registrant contact. 2303 OK
- creation of a domain with a nonexistent billing contact. 2303 OK
- creation of a domain with a nonexistent tech contact. 2303 OK
- creation of a domain with a nonexistent onsite contact. 2303 OK
- creation of a domain that is just created. 2306 IN USE OK
- creation of a domain that is already applied for. 2306 IN USE OK

Check domain

- Try to create and check for a domain without .eu. 1000 OK
- Try to create and check for a domain with .eu. 1000 OK
- Try to create and check for 3 domain. 1000 OK
- **Check for a .it domain. 1000 ! It's that OK!? a part of response:**
`<domain:name avail="1">iit-1127902894842.it.eu</domain:name>`
- Check for an empty domain. 2001 OK

Delete domain

- create and delete a domain. 1000 OK
- delete a domain that is already applied for. 2304 Object status prohibits operation OK
- delete a nonexistent domain. 2304 OK
- **set a past delete date. 1000 !It's possible to send a past delete date! It's OK?**
- set a date in another format. 2001 OK
- set an empty date. 2001 OK
- delete a domain without delete date. 2001 OK

Info domain

- Try to create and info for a domain without .eu. 1000 OK
- Try to apply and info for a domain without .eu. 1000 OK
- Try to create and info for a domain with .eu. 1000 OK
- Info for a .it domain. 2303 OK
- Info for an empty domain. 2001 OK

Update domain

- Try to add and remove an onsite contact in the same update.
2308 Data management policy violation OK
- Try to update a domain with a name server under .eu specifying the ip address.
2306 Parameter value policy error OK
- Try to create and update a domain adding an onsite contact. 1000 OK
- Try to update a domain after an apply. 2304 Object status prohibits operation OK
- Try to create and update a domain adding a non-existing contact. 2303 Object does not exist OK
- Try to create and update a domain adding an empty contact. 2001 OK
- **Try to send an update with make no changes. 1000 ! It's OK ?**
- Try to update a domain with an empty domain name. 2001 OK
- Try to update a domain with a nonexistent domain name. 2304 OK

Undelete domain

- Try to create, delete and undelete a domain. 1000 OK
- Try to create and undelete a domain. 2304 "Domain ... not scheduled for delete" OK
- Try to undelete a nonexistent domain. 2304 OK

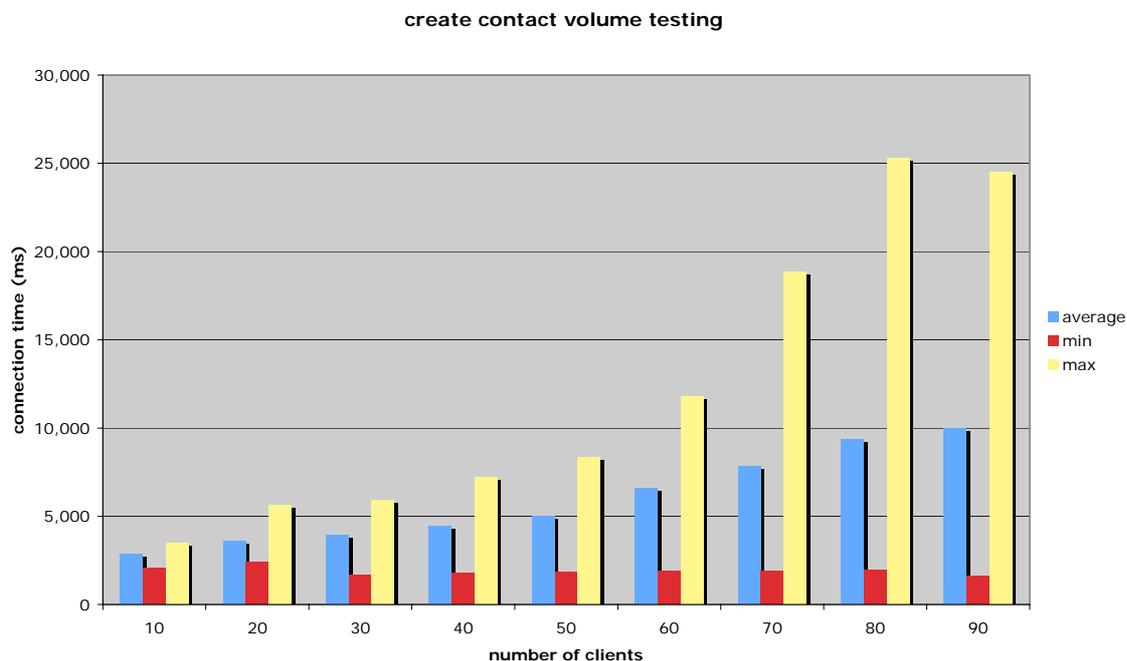
Preliminary Workload Testing

Although intensive workload stress testing of the final system is not possible at the moment, we carried out some tests on the old system in order to understand the trend of growth of the time observed by the user when increasing the server load.

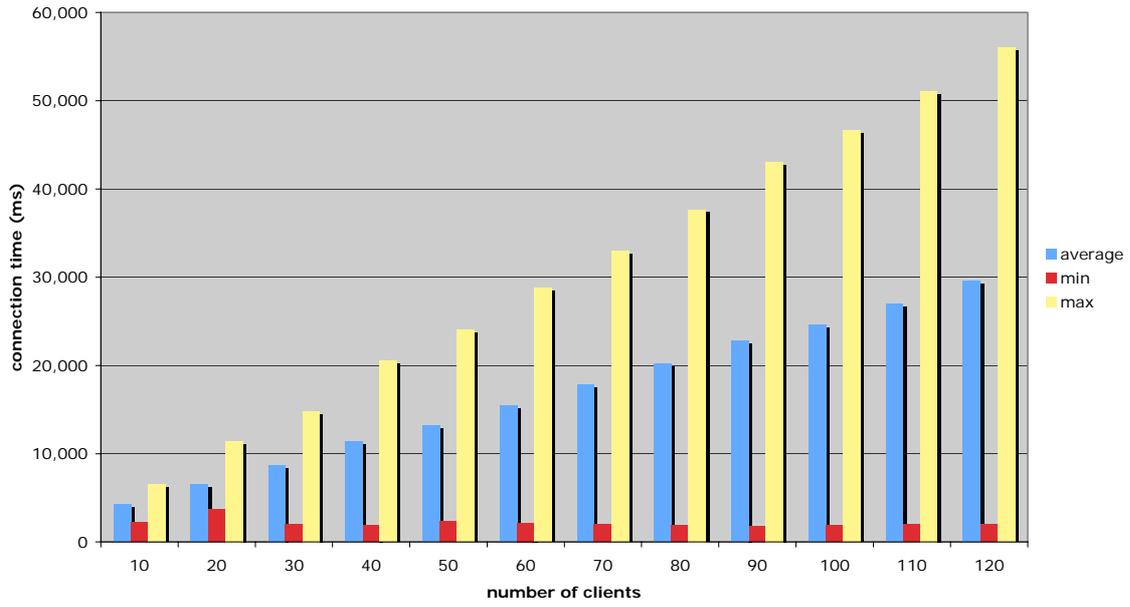
We focused on 3 commands: create contact, apply domain and create domain. We performed the tests launching load generators from 2 hosts located at CNR, Pisa, Italy. Each generator instantiated a custom number of threads in parallel (increased by 5 for each test) and then connected them to the server at the same time. We synchronized the generators and let each of them launch the same number of clients at the same time.

The times we measured refer to the interval between the instant when the client sends its SYN packet to connect, to when it closes the socket. The test actually stopped when the client received a reply from the server after creating or applying for a domain, since the server forces the connection to close after those commands.

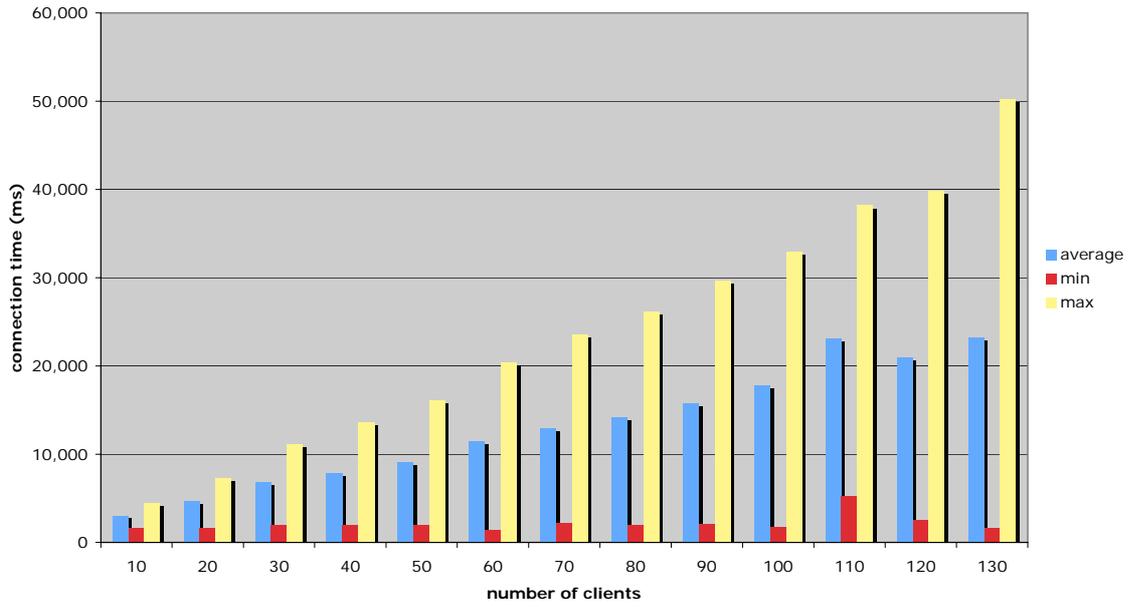
After the tests we calculated minimum, maximum and average values on the results of both hosts grouped by starting time. The results presented in the graphs below show the total connection time perceived by clients, which includes SSL handshake, login, and the command to test. In the case of contact creation this time includes the logout command and its reply as well. By the end of the create contact and the create domain test (both performed on Monday, Sept. 26th) we had some clients (4 of 100 and 2 of 140 respectively) experiencing a remote closing of the connection during SSL handshake.



apply domain volume testing



create domain volume testing



Workload Stress Testing

Lab set up and test description

In this set of experiments we highlighted the average connection time perceived by an increasing number of clients. We focused on 2 commands: create contact and apply domain (as required by EURid partners).

We set up at CNR, Pisa, Italy a laboratory with 20 homogeneous hosts. The main characteristics of the hosts were:

- CPU: Intel Pentium 4 CPU 1.60GHz, 256 KB cache
- RAM: 256 MB
- Ethernet controller: Intel Corp. 82801BA/BAM/CA/CAM
- OS: Linux 2.6.9, Java 1.5.0

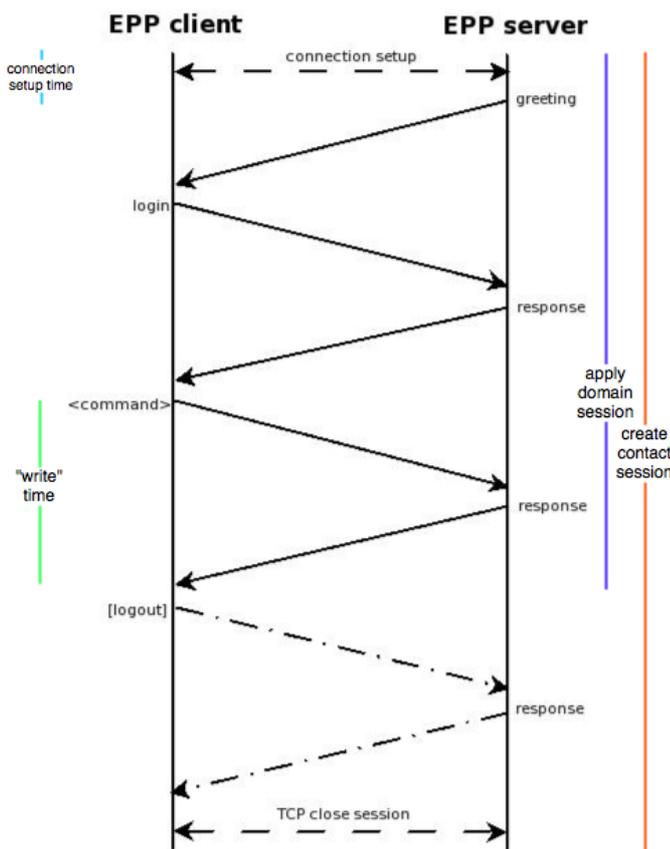
We performed the tests launching synthetic traffic generators distributed on 20 these different hosts. Each generator instantiated a custom number of threads in parallel and then connected them to the server **at the same time**. We started with 20 clients (one client on each host) and increased by twenty (again, by one on each host) at each step of the test (i.e. 20, 40, 60, ...).

We synchronized the generators and let each of them launch the same number of clients and connect them at the same time: the highest delay that we measured between the first and the last of 50 SYN packets sent from one of our hosts was about 35 milliseconds. Considering also the traceroute from us to the server (shown in the figure below), we can assume that:

- our traffic generator guarantees scalability to 50 clients (1000 combined clients)
- the route from us to the server is not a bottleneck

```
traceroute to 212.190.251.50 (212.190.251.50), 30 hops max, 38 byte packets
 1 aule-static26 (146.48.89.254) 0.440 ms 0.378 ms 0.353 ms
 2 rtg-cnrpi.pi.garr.net (193.206.136.37) 0.427 ms 0.442 ms 0.475 ms
 3 rt-pi1-rt-to1.to1.garr.net (193.206.134.73) 4.764 ms 4.748 ms 4.723 ms
 4 rt-to1-rt-mi2.mi2.garr.net (193.206.134.41) 7.470 ms 7.564 ms 7.425 ms
 5 213.248.71.161 (213.248.71.161) 7.429 ms 7.654 ms 7.358 ms
 6 * * *
 7 * * *
 8 POS3-1.BR2.FFT1.ALTER.NET (146.188.67.85) 19.581 ms 19.540 ms 19.916 ms
 9 so-0-2-0.TR1.FFT1.ALTER.NET (146.188.6.109) 19.998 ms 19.747 ms 19.643 ms
10 so-6-0-0.TR2.BRU2.ALTER.NET (146.188.9.46) 33.459 ms 33.987 ms 33.478 ms
11 so-6-0-0.XR1.BRU2.ALTER.NET (146.188.9.150) 33.665 ms 33.497 ms 33.462 ms
12 POS1-0.GW3.BRU2.ALTER.NET (146.188.2.82) 34.573 ms 33.470 ms 34.464 ms
13 uu194-7-12-146.unknown.uunet.be (194.7.12.146) 33.404 ms 33.756 ms 33.267 ms
14 uu194-7-12-146.unknown.uunet.be (194.7.12.146) 34.217 ms 33.692 ms 33.508 ms
15 eurid.customers.be.uu.net (212.190.251.50) 35.175 ms 35.508 ms 34.981 ms
```

The below figure briefly shows the composition of an EPP transaction and the times that we measured during the tests.

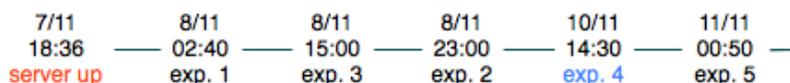


We measured three values for each parallel transaction:

- the *total* connection time (shown in the figure as “apply domain session” and “create contact session”), which refers to the interval between the instant when the client sends its SYN packet to connect, to when it closes the socket;
- the *connection setup* time, which is the time measured between the SYN packet sent by our client and the end of the SSL handshake;
- the “*write*” time as perceived by the client, which refers to the time between when the client sends the apply domain request and when it receives the server’s reply.

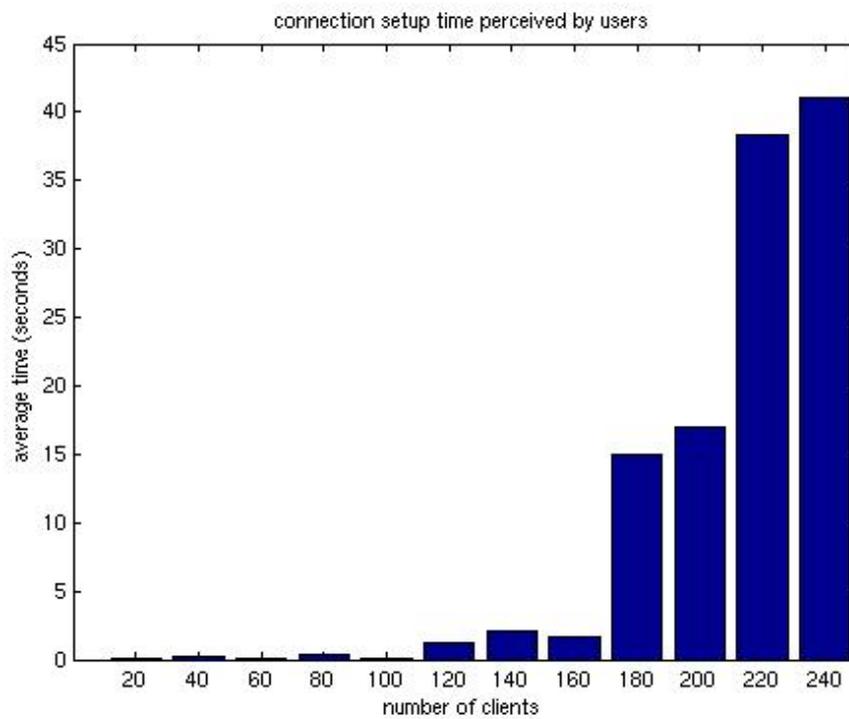
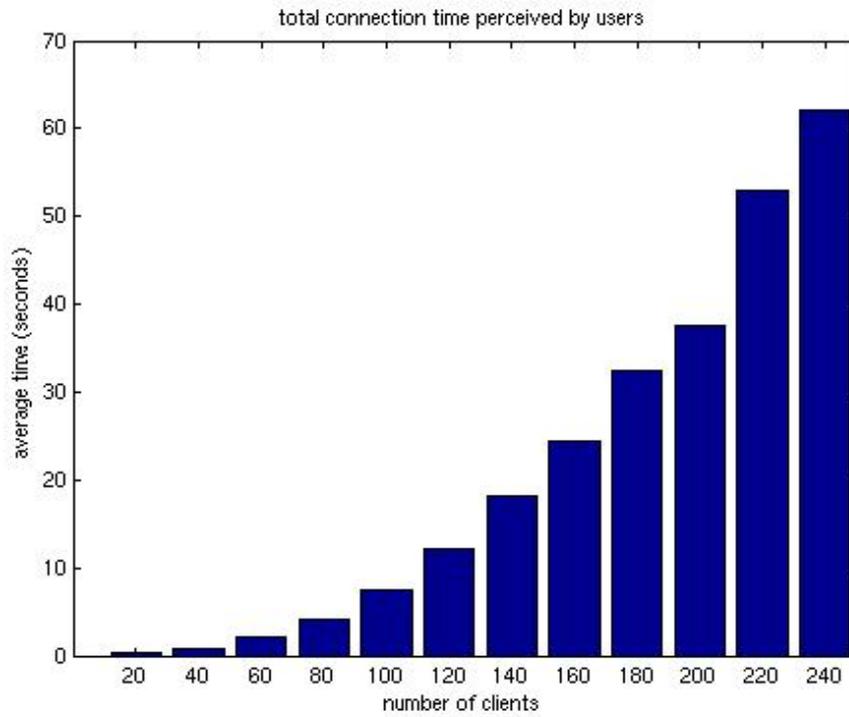
We reported the average value of each set of measured values, except for the “write time”: once the connection is established, the latency on the database perceived by our clients seemed to be constantly around 500 ms, so not critical at all compared to the total connection time as we will show. For what concerns the connection setup time, which appears to be very critical, we reported a brief statistical analysis in terms of percentiles as well, which seems to show that the peaks shown in the graphs are due to the activity of an increasing number of clients which experience very long delays in the connection setup. We reported the results starting from 20 clients (1 client on each host) and ending with the last consistent experiment, which we considered the last experiment with a percentage of failed transactions around 5 to 10% depending on the experiment.

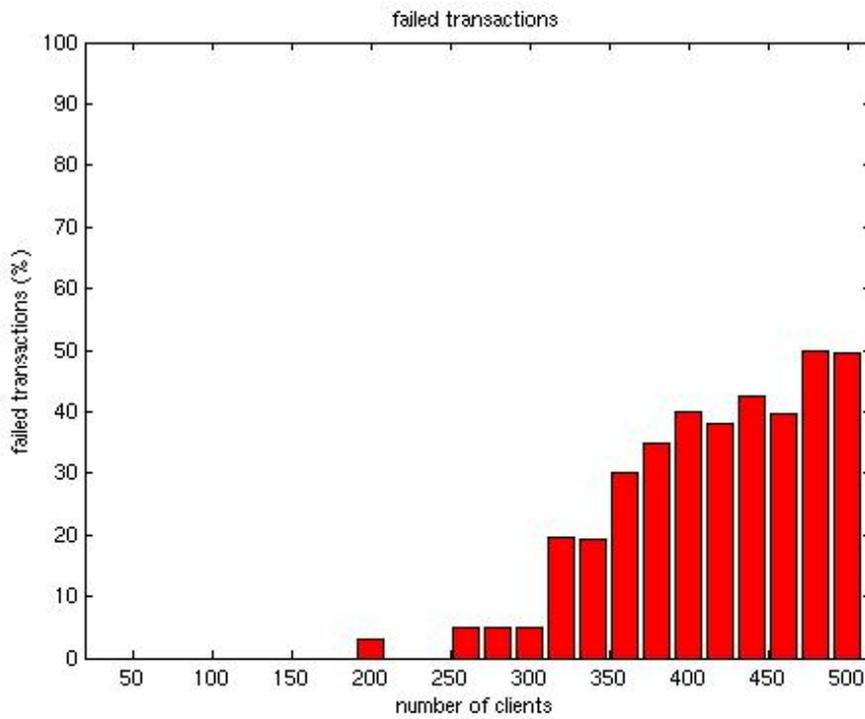
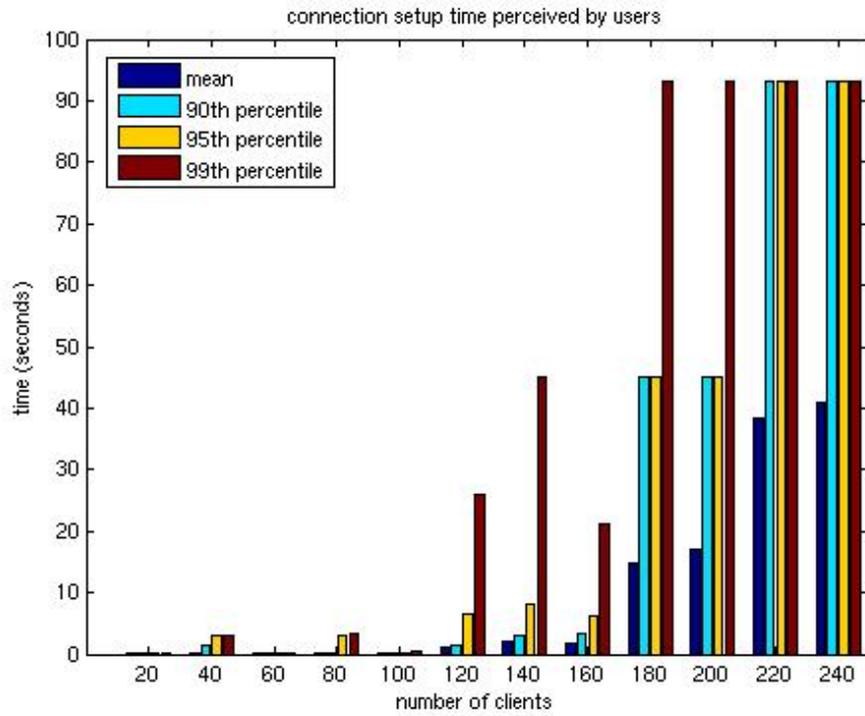
We report here the results of 5 experiments: each step of each test ran 10 minutes after the beginning of the previous step to guarantee the absence of overlays, so a test involving up to 1000 clients lasted about 8 hours. Here is a timeline of the starts of the experiments that follow, all performed in the week from 7/11/2005 to 11/11/2005.



Experiment 1 – Create Contact

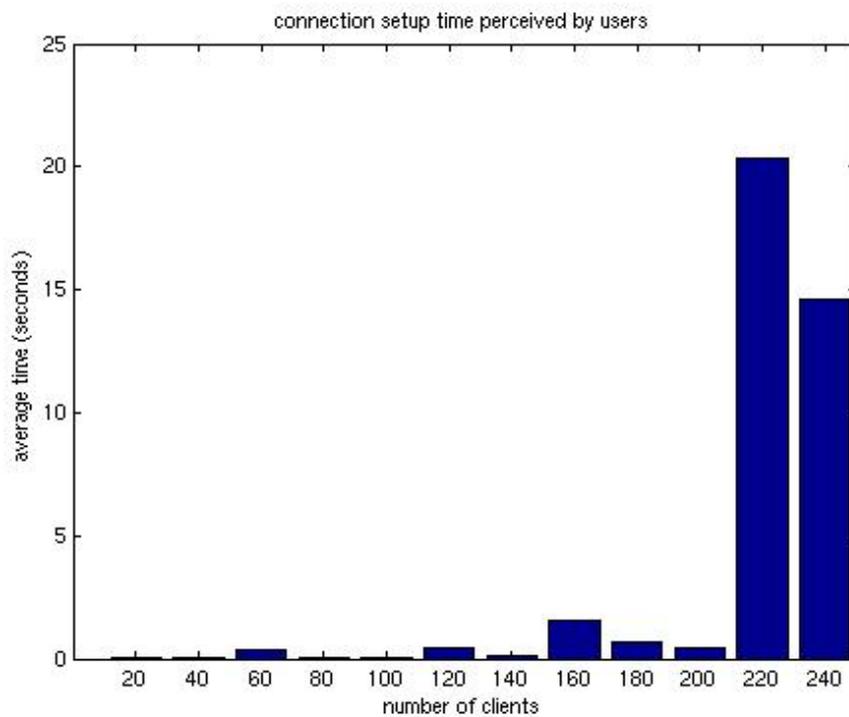
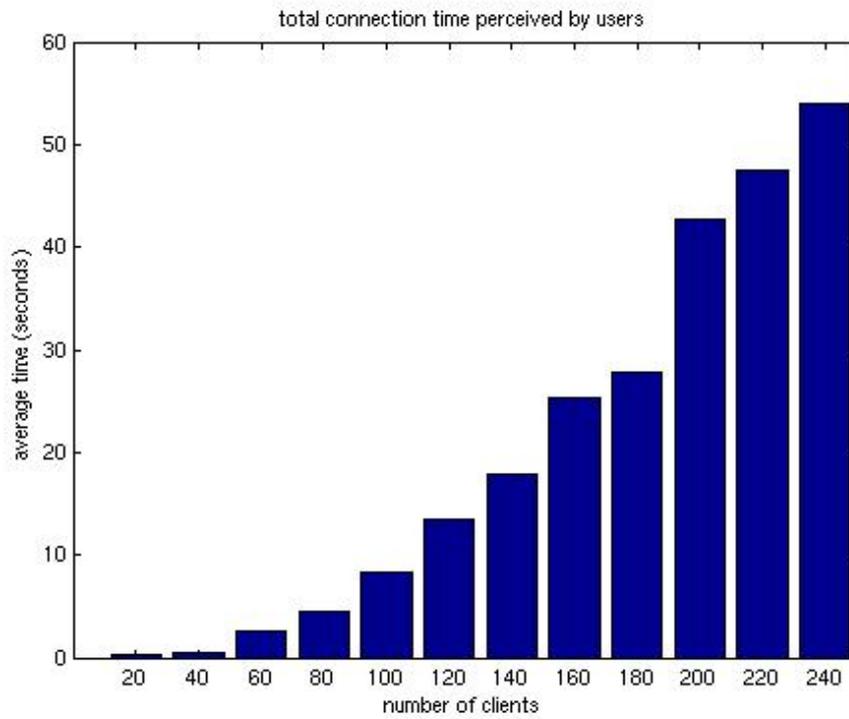
20 to 500 clients (1 to 25 per host) performing a create contact using the same registrar ID (t002000).

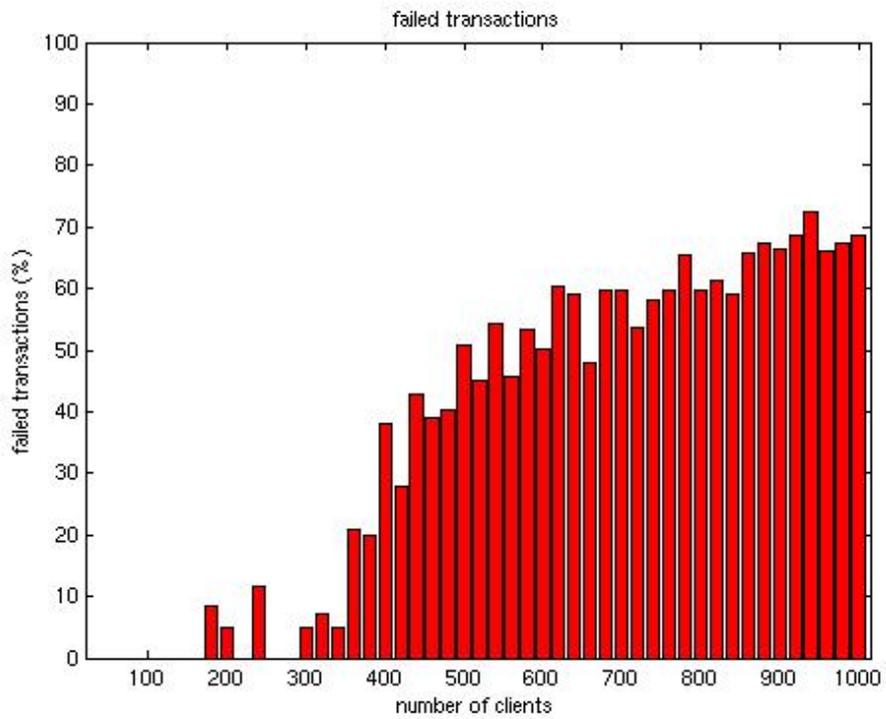
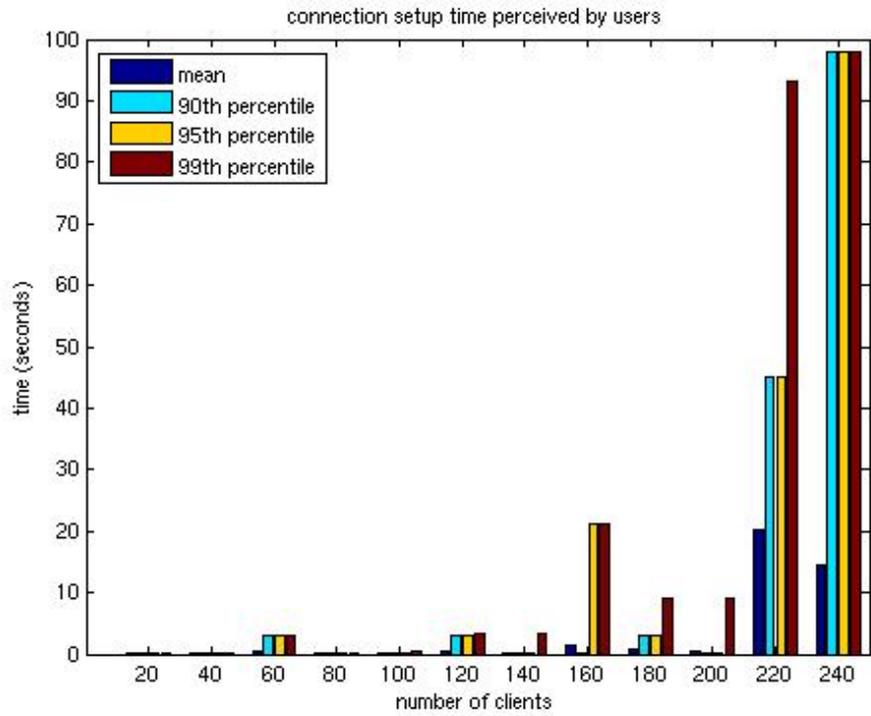




Experiment 2 – Create Contact

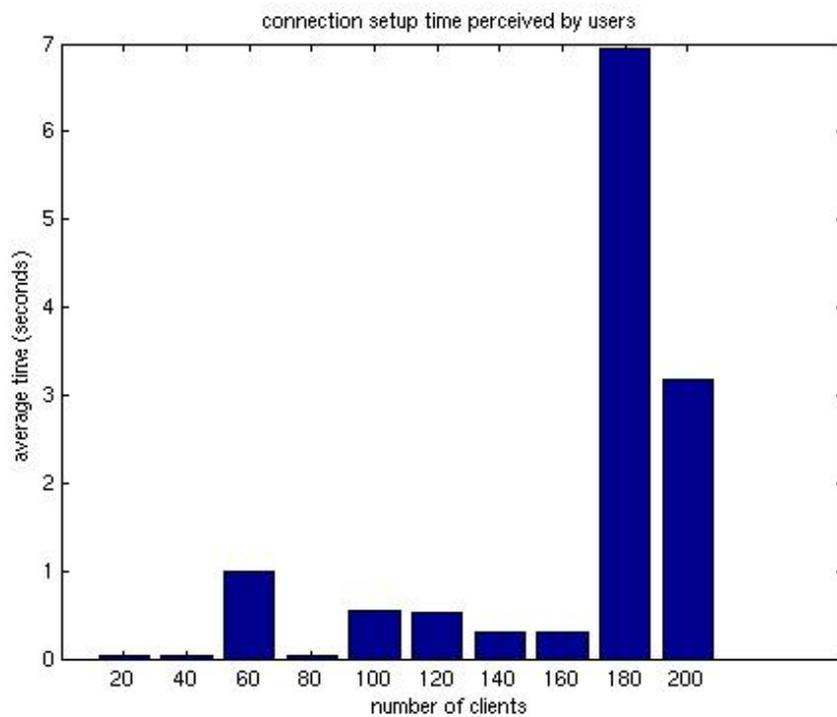
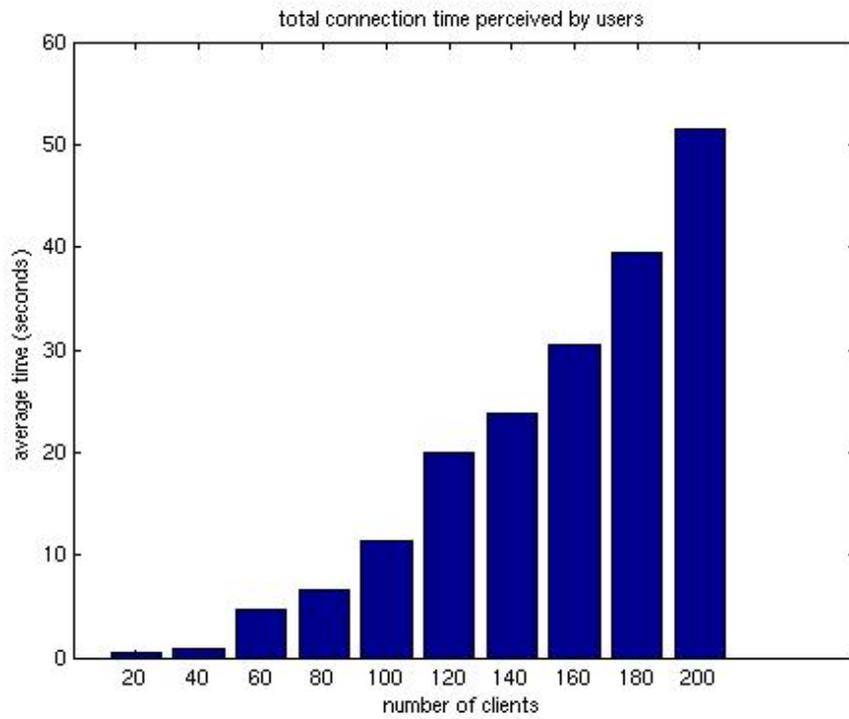
20 to 1000 clients performing a create contact using the same registrar ID (t002000).

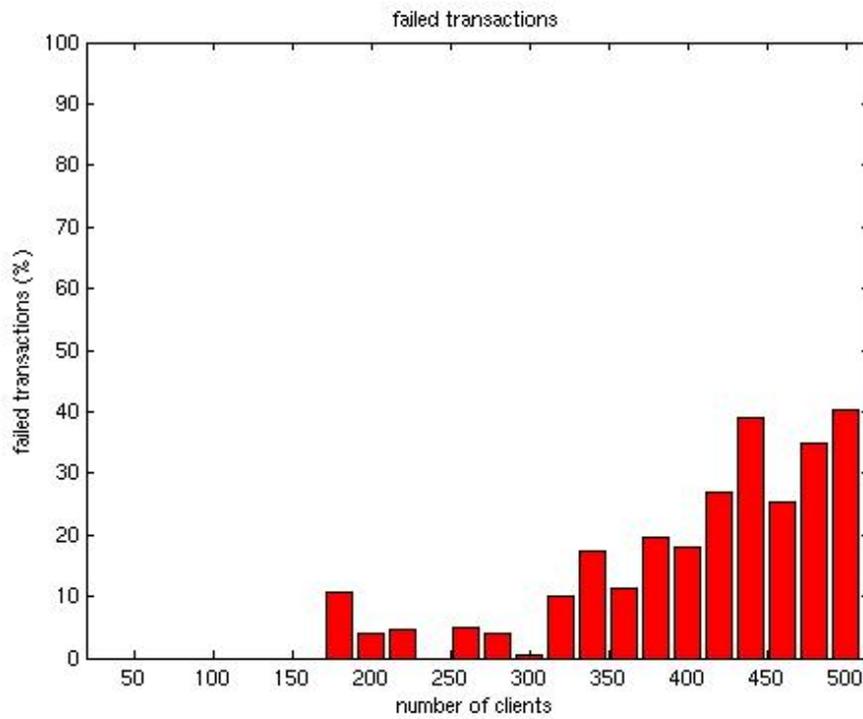
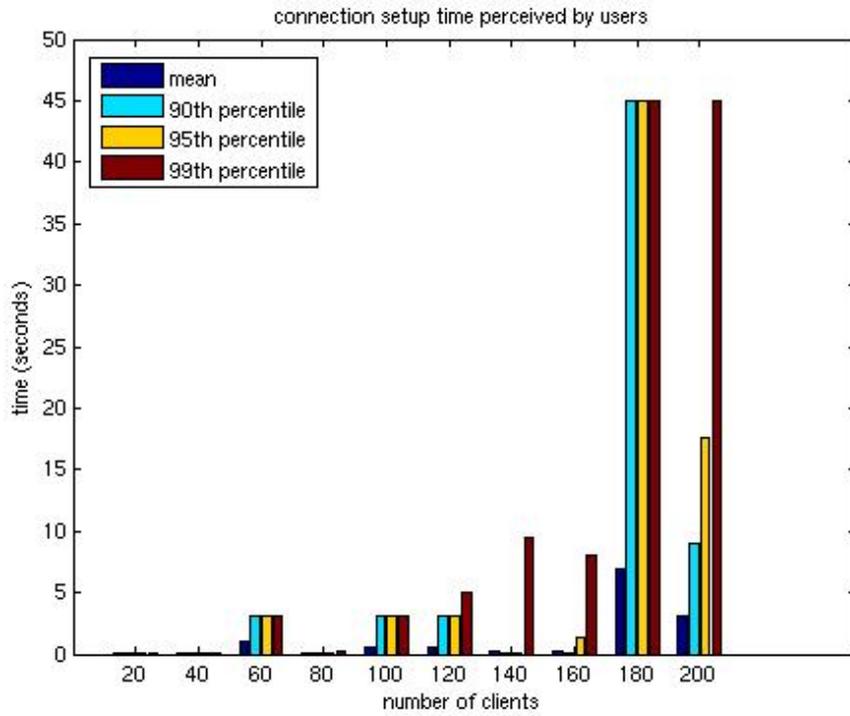




Experiment 3 – Apply Domain

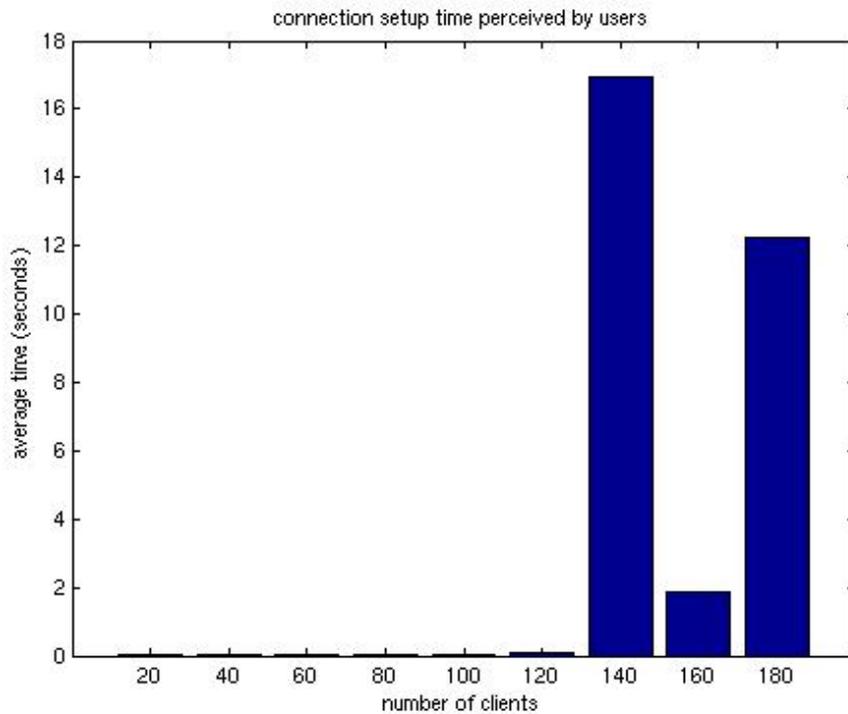
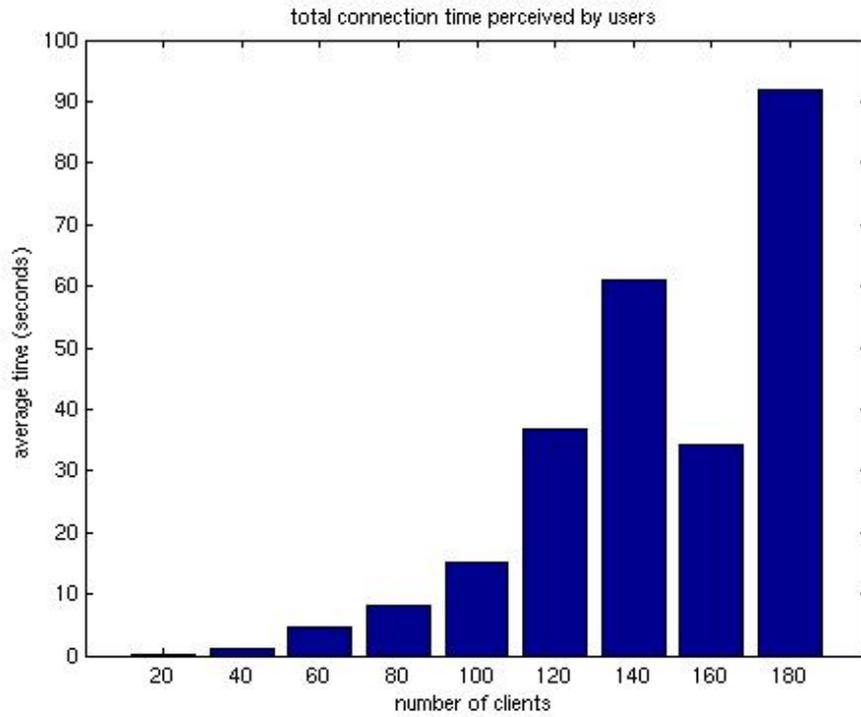
20 to 500 clients performing an apply domain using the same registrar ID (t002000).

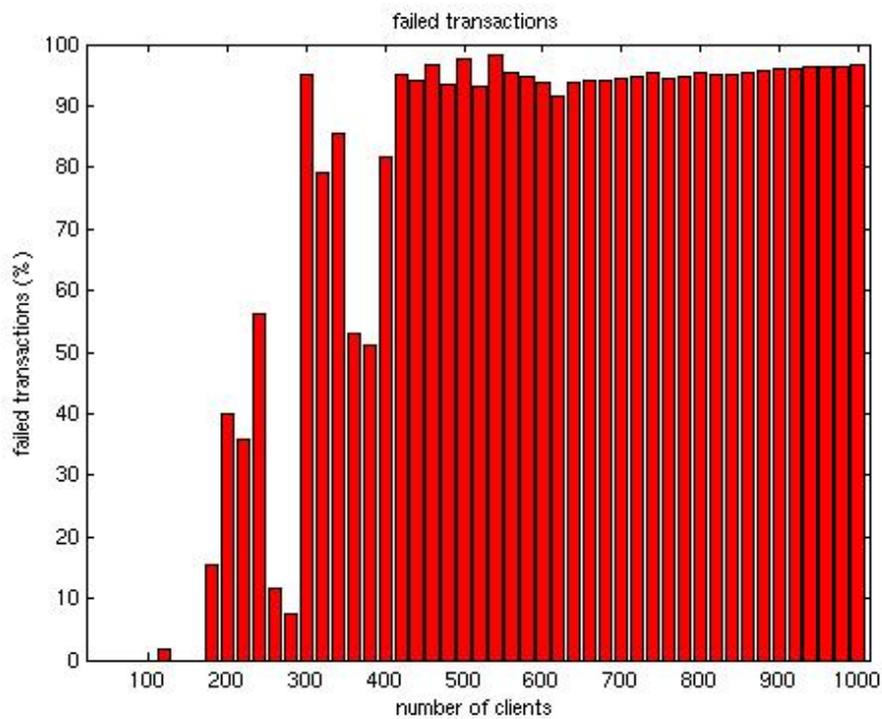
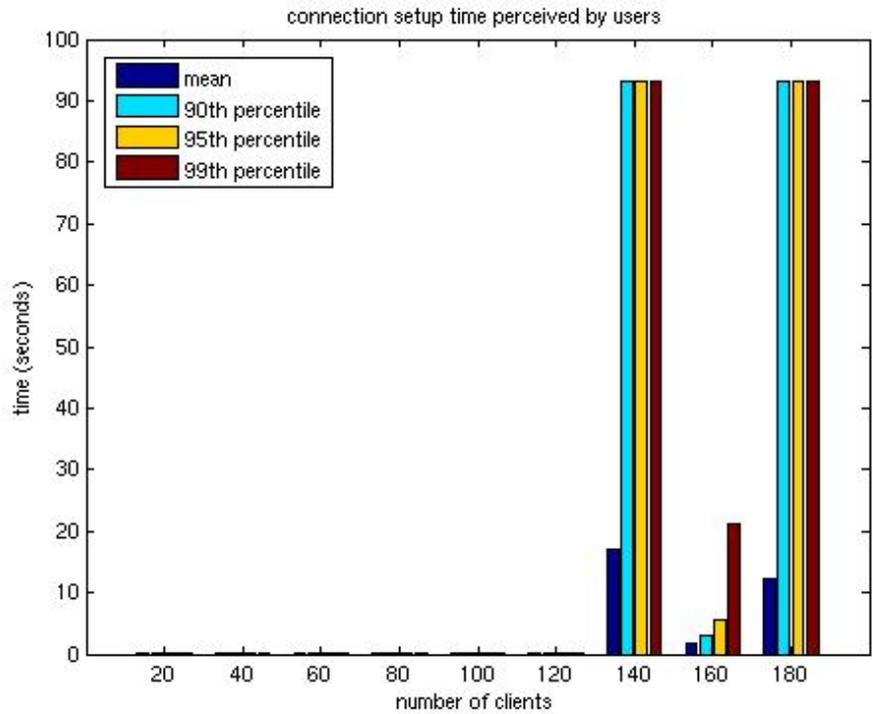




Experiment 4 – Apply Domain

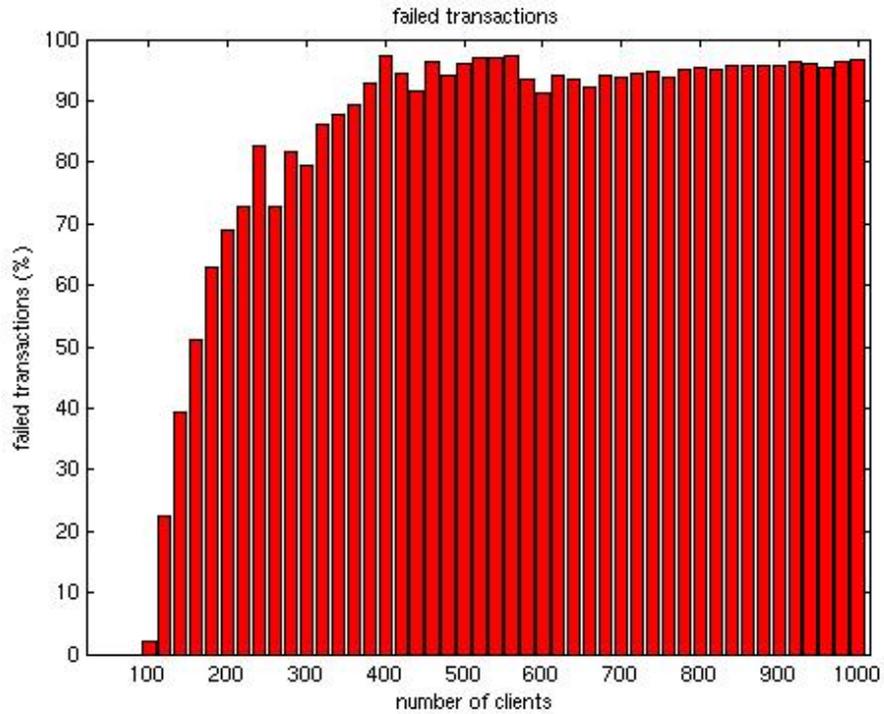
20 to 1000 clients performing an apply domain using different registrar IDs (t002000 to t002999).





Experiment 5 – Apply Domain

20 to 1000 clients performing an apply domain using different registrar IDs (t002000 to t002999). This is actually a repetition of *experiment 4* on Fri Nov 11 00:43:00 CET getting much lower performances.



However a possible cause of this even lower performance versus the previous experiments maybe due to a contemporary workload test carried out locally from our Belgium partner. Tests revealed 3 kinds of errors whose frequency is detailed in Appendix A. The dominant error is connection time-out, which starts to occur after a certain number of connections.

Final test

On the basis on acquired test results and monitoring server-side logs, our partner dns.be was able to fix a bug in the connection set-up code.

We therefore ran a final test on Friday November 25th were all the partners and some registrar were involved, generating load and DoS attacks. We used our 20 hosts just to measure what QoS was still guaranteed in terms of successful transactions.

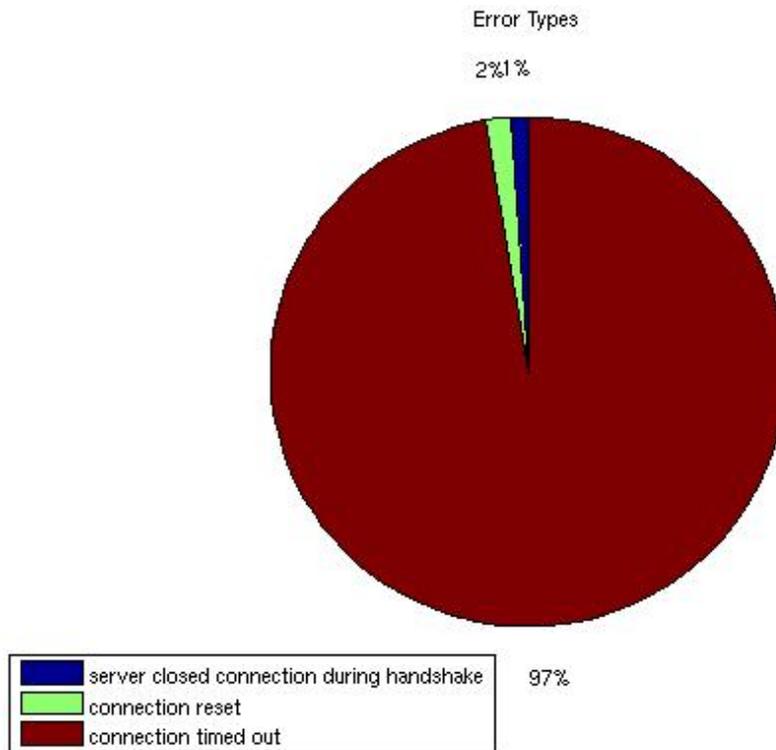
From the results presented in the table below we can conclude that under the load generated the server responded very well and with moderate fairness.

We had two types of errors due to unrecognisable responses from the server, indicated as Number Format and EOF. The SSL handshake errors are actually the same as the other tests: the server closes the connection during the handshake. Note that the number of clients run differs from host to host because of the different number of errors.

	tests run	successful	errors	Error types		
				Number	EOF	SSL
host # 1	2083	2061	22	12	5	5
host # 2	2074	2057	17	9	0	8
host # 3	2077	2054	23	13	3	7
host # 4	2086	2066	20	13	0	7
host # 5	2095	2064	31	25	0	6
host # 6	2100	2076	24	18	2	4
host # 7	2093	2068	25	18	0	7
host # 8	2097	2073	24	15	3	6
host # 9	2097	2072	25	15	2	8
host # 10	2097	2069	28	20	2	6
host # 11	2099	2068	31	22	3	6
host # 12	2098	2071	27	19	3	5
host # 13	2083	2065	18	9	0	9
host # 14	2094	2073	21	14	1	6
host # 15	2095	2062	33	26	2	5
host # 16	2099	2078	21	15	1	5
host # 17	2094	2062	32	24	1	7
host # 18	2088	2061	27	19	1	7
host # 19	2071	2053	18	9	3	6
host # 20	2112	2086	26	19	3	4

Appendix A: error types

We encountered 3 kinds of errors whose frequency is shown in the below graph.



Server closed connection during handshake

This error seems to happen frequently when around 200 clients try to connect. You can see that zone clearly in the failure graphs previously presented. It might be due to a late TCP retransmission.

Connection reset

This error doesn't seem to happen in a well-defined condition: what happens is that the server just sends a RST packet. Like the previous, it might be due to retransmission.

Connection timed out

This is the dominant error because it starts to happen after a certain number of connections. It can be seen clearly in the failure graphs (on the right from certain point to the end of the test). What we have in our connection dumps is just all the SYN packets of the clients that we instantiated and no response from the server.

Appendix B: developed software

EuTester: package class diagram

