

TECH NOTES

JABBER XCP PERFORMANCE TESTS — FROM 100,000 TO ONE MILLION CONCURRENT USERS

The Jabber Extensible Communications Platform™ (Jabber XCP™) is the most scalable, extensible, carrier-grade, multi-protocol presence and messaging platform available on the market today. It is uniquely capable of powering an unlimited variety of presence-enabled, real-time applications, services, devices, and systems.

The intention of this document is to assist information technology (IT) professionals in understanding the unparalleled performance capabilities of a Jabber XCP solution by describing how the servers perform under load and scale with additional resources, and by explaining how Jabber, Inc. tests its configurations.

EXECUTIVE SUMMARY

In separate tests, Jabber, Inc. partnered with France Telecom Research and Development and Sun Microsystems, Inc. to showcase Jabber XCP's ability to support 100,000 concurrent users and, subsequently, to scale to more than one million concurrent users without reaching capacity and hindering user performance.



FRANCE TELECOM RESEARCH AND DEVELOPMENT

Jabber, Inc. conducted a 48-hour performance test with France Telecom Research and Development, proving that the Jabber XCP platform successfully handles at least two million subscribers and 100,000 concurrent users, while running on a single Jabber XCP router (jabberd) with a pool of Connection Managers. The test focused on the jabberd (message routing and sessions), Connection Manager (client/server connection control), and xdb_oracle (user information library) components of Jabber XCP. Results of the test indicated that CPU utilization never exceeded 73%, memory usage was below 895 MB, and average message latency was below 0.29 seconds.

SUN MICROSYSTEMS, INC.

As a partner in the Sun Partner Advantage Program for Independent Software Vendors (ISVs), Jabber, Inc. utilized Sun's Benchmark Center labs in Edinburgh, Scotland to perform load and scalability testing on the Jabber XCP platform. Jabber, Inc. employed a phased approach to achieve the objectives of the test, ensuring that the platform could successfully provide highly scalable, presence-oriented services. The purpose of the first phase was to load and support 420,000 concurrent users, and served as a baseline for the second phase of scalability testing. The second phase focused on incremental hardware additions to accommodate over one million concurrent users, thereby demonstrating the massive scalability of the environment.

**JABBER XCP
CAN MEET THE
PERFORMANCE
NEEDS OF
ORGANIZATIONS
OF VIRTUALLY
ANY SIZE**

Phase I utilized a load scenario that averaged requests and data collected from customers and prospects, and proved that Jabber XCP successfully supports 420,000 concurrent users on two Sun Fire™ T2000 servers. Results indicated that CPU utilization did not exceed 48% and memory usage stabilized at 3.3 GB with an average authentication latency of approximately 0.06 seconds.

In Phase II, Jabber, Inc. deployed six more Sun Fire T2000 servers to support the substantial increase in users. The test demonstrated linear scalability with more than one million concurrent users. Jabber XCP exceeded performance expectations even with this extremely high level of users, as CPU utilization never exceeded 44%, memory usage stabilized at 2.1 GB, and authentication latency was approximately 0.05 seconds. As the systems were not fully utilized, more users could have been supported in this environment. Additional servers could have been added to meet scalability requirements beyond the scope of the tests.



THE TEST RESULTS ILLUSTRATE THE SCALABILITY AND RELIABILITY OF JABBER XCP

The results of these tests clearly illustrate the maturity, scalability, and reliability of the Jabber XCP platform, and prove that Jabber XCP addresses the scalability needs for massive presence and messaging deployments. To Jabber, Inc.'s knowledge, no other presence and messaging platform has proven this level of performance and linear scaling.

JABBER, INC. RELIABILITY AND PERFORMANCE TESTING

Reliability and performance testing are key components of the Jabber, Inc. quality assurance (QA) process. Jabber XCP consists of a distributed architecture of software components, and each component is evaluated in a performance test. The Jabber, Inc. QA team uses a proprietary load simulator that mimics the behavior of multiple users connected to a Jabber XCP server with probability-based, event generation.

The simulator sends Extensible Messaging and Presence Protocol (XMPP) stanzas that comply with Internet Engineering Task Force (IETF) Requests For Comments (RFC) 3920 and 3921. XMPP "events" (e.g., send chat message, change presence, and rename roster group) are defined in the simulator and occur in percentages based on Jabber, Inc.'s experience with service providers and enterprise customers. The simulator runs a second simulation, which logs in two known users and tracks the latency of messages sent between them through the system. It captures CPU, memory, network, and system load performance data while recording summary data at varying intervals throughout the duration of the tests.

100,000 CONCURRENT USERS TEST

METHODOLOGY AND TEST MODES

To accomplish this test, QA engineers configured four servers as simulators, three servers as Connection Managers (each running two Connection Manager instances), one as the server for the Jabber XCP router (jabberd), and a final server was used for the database. The test was configured on a gigabit network. Tests included:

STANDARD PROBABILITY-BASED TEST

This is the default configuration for the simulator. Probabilities are assigned to each XMPP event type simulated. (Probabilities are the percentage of time that the XMPP event will occur.) The order in which events occur is random.

MONITOR TEST

This is an order-based test. Two unique users are created for the monitoring test. Messages are sent between these two users during the load test. Jabber, Inc. records the time difference between when an event is triggered and when a response is received from the server.

USER LOAD AND PERFORMANCE TEST RESULTS

In this test, the Jabber XCP server managed a load of two million registered users (each with a roster of 20 contacts), 100,000 concurrent users logging in at the rate of 120 per second, and a 75-second delay between message, presence, vCard lookup, and subscription activities per logged-in user. The test examined three major components:

- **Jabber XCP Router (jabberd)** – Controls message routing and sessions
- **Connection Manager** – Controls connections between client and server
- **xdb_oracle** – A library loaded into jabberd that handles storage of user information (e.g., registration, authentication, roster information, vCards, and offline messages)

Events were triggered over a 48-hour period to test login, authorization, messages, roster changes, presence, and vCard get/set. Jabber, Inc. recorded memory and CPU loads during the initial, login, and final stages. The initial stage measured the server state before the simulation started. The login stage measured the server state after all the concurrent users logged in. The final stage measured the server state after the simulation completed all events. The results are summarized in the following tables:

Memory Usage

Component	Initial	Login	Final
Jabber XCP Router (jabberd)	60928 KB	509952 KB	916480 KB
Connection Manager 1	3876 KB	182272 KB	200704 KB
Connection Manager 2	3876 KB	184320 KB	203776 KB
Connection Manager 3	143360 KB*	323584 KB*	344064 KB*
Connection Manager 4	4136 KB	183296 KB	201728 KB
Connection Manager 5	3872 KB	183296 KB	204800 KB
Connection Manager 6	143360 KB*	322560 KB*	344064 KB*
Oracle Server	10592 KB	537600 KB	1050624 KB

*Problems with memory reporting on the version of the UNIX command 'top' used on AS 4.0 caused initial memory to report as 140MB

100,000 USERS TEST SUMMARY RESULTS

- Maximum CPU Usage: 72%
- Maximum Memory Usage: 895 MB (jabberd)
- Average Message Latency: 0.28288 Seconds
- Average Presence Latency: 0.35150 Seconds

CPU Usage

Component	Initial	Login	Final
Jabber XCP Router (jabberd)	0	1.9%	47.5%
Connection Manager 1	0	1.1%	0.5%
Connection Manager 2	0	1.9%	0.2%
Connection Manager 3	N/A*	N/A*	N/A*
Connection Manager 4	0	2.5%	0.3%
Connection Manager 5	0	2.1%	0.3%
Connection Manager 6	N/A*	N/A*	N/A*
Oracle Server	0	0.3%	34.3%

*Problems with CPU Reporting on the version of the UNIX command 'top' used on AS 4.0 make this data unreliable

MEMORY FOOTPRINT PER USER

Two important components to consider when evaluating memory footprint per user are the Jabber XCP router (jabberd) and Connection Manager.

JABBER XCP ROUTER (JABBERD)

The initial memory for the server was 60928 KB, while the memory after login and activity from 100,000 users was 916480 KB. Dividing the difference by 100,000 yields a per-user footprint of approximately 8.555 KB.

Note that memory footprint depends on many factors, such as size of roster, type of authentication, and so forth. Different usage patterns may result in different outcomes.

CONNECTION MANAGER

The highest memory for the Connection Manager was 204800 KB, minus its starting value of 3876 KB, multiplied by six Connection Managers and divided by 100,000 users gives us approximately 12 KB. Memory footprint depends on many factors, such as size of roster, Secure Sockets Layer (SSL) usage, type of authentication, and so forth. Different usage patterns may result in different outcomes. Also, note that the memory levels after login (for both jabberd and Connection Manager) will grow for a period of time as the system adjusts to the load while performing tasks such as caching data. Therefore, all of the calculations are based on the state of the system at the end of the 48-hour period.

100,000 USERS TEST ENVIRONMENT

To handle 100,000 concurrent users and two million registered users, the test used a pool of three servers each running two Connection Managers, one server running the Jabber XCP router (jabberd) and one server running Oracle.

USER ENVIRONMENT

- 100,000 concurrent users
- 2 million registered users

SOFTWARE ENVIRONMENT

- Red Hat AS 3.0 (Jabber XCP 4.2)
- Red Hat AS 3.0
(Connection Managers 1, 2, 4, 5)
- Red Hat AS 4.0
(Connection Managers 3, 6)
- Oracle 9i (9.2.0.6)

HARDWARE CONFIGURATION

The tests used 5 identical servers
(1 Jabber XCP, 3 Connection Managers,
1 Oracle database)

- 4 Intel Xeon MP 2.8GHz Hyperthreaded with 512 L2 Cache
- Single 35 GB 10K Ultra 320 SCSI
- Adaptec AIC7902 Ultra 320 SCSI adapter
- 4 GB 100MHz DDR RAM
- Intel Pro/1000 NIC

STORAGE AREA NETWORK (SAN) INFRASTRUCTURE

- Nexan SATABlade 3.2 TB Storage
- QLogic Fibre HBA QL-2340 (installed in Oracle server)

NETWORK INFRASTRUCTURE

- Gigabit Ethernet network
- 3COM 3824 24-Port GB switch

TEST PARAMETERS

- 48-hour duration
- 120 logins per second
- 20 contacts per user
- 200 byte message size
- 75 seconds between each message activity per user

XMPP EVENTS TRIGGERED

- Connection
- Registration
- Login
- Presence changes
- Send message to online users
- Roster add, delete, get, rename group
- vCard get/set

ONE MILLION CONCURRENT USERS TEST

METHODOLOGY AND LOAD TEST CONFIGURATION

To accomplish this test, QA engineers configured 16 servers as simulators, two servers as Connection Managers, six Jabber XCP router servers (jabberd), and one database server. The servers were connected to two private gigabit Ethernet networks.

The following test scenario was used to configure load test simulation:

- The number of users in the database was 1,800,000, and the number of active users was between 420,000 and 1,020,000
- The average roster size was 30, and the simulated number of active users per roster was seven
- Each user performed an action every 120 seconds
- The probability that an action was to send a message was 83.5%, the probability that an action was to log off and immediately log back in was 2.5%, the probability that the action was a series of roster operations (including add, rename, and remove) was 0.5%, and the probability of a presence change was 13.5%

ONE MILLION USERS TEST SUMMARY RESULTS

- Maximum CPU Usage: 44%
- Maximum Memory Usage: 2.1 GB
- Highest Average Message Latency: 0.03561 Seconds
- Highest Average Presence Latency: 0.03957 Seconds

USER LOAD AND PERFORMANCE TEST RESULTS

CPU USAGE AND MEMORY

The CPU usage and memory peak was seen immediately following login and was due to the high number of logins per second during the beginning of the test when all users logged in to the system. The peak period lasted for approximately 30 minutes, after which steady state was achieved and continued for the remainder of the tests. CPU utilization never exceeded 44%, and memory usage never exceeded 2.1GB. The following chart shows steady state memory and CPU usage for both the router (session layer) and Connection Manager (client connection layer).

Beginning the test with 420,000 concurrent users, CPU utilization was 48% and memory usage was 3.3 GB. By the time the test scaled to 1.02 million concurrent users, CPU utilization was 44% and memory usage was stabilized at 2.1 GB.

CPU Usage and Memory

Concurrent Users	Router Server Machines	Router Memory Per Server Machine	Router CPU Per Server Machine	Connection Manager Server Machine/ Cores Per Server Machine	Connection Manager Memory Per Server Machine	Connection Manager CPU Per Server Machine
420,000	1	3.3 GB	48%	1 / 4	6.4 GB	52.8%
420,000	2	2.3 GB	58%	1 / 4	6.4 GB	52.8%
576,000	3	2.7 GB	61%	1 / 5	8.2 GB	80%
720,000	4	2.2 GB	57%	1 / 6	10.0 GB	70%
870,000	5	2.1 GB	57%	2 / 4	6.2 GB	40%
1,020,000	6	2.1 GB	44%	2 / 5	7.2 GB	55.2%

ORACLE SERVER

Even with more than one million concurrent users, the Oracle server used less than two of the 32 virtual processors available to it during steady state and less than four of the 32 available during peak load.

JABBER XCP ADDRESSES THE SCALABILITY NEEDS FOR MASSIVE DEPLOYMENTS

AVERAGE LATENCIES

True end-to-end latency was determined by inserting a precise time-stamp in each XMPP stanza that was sent, and then calculating the difference when the stanzas were received. Four separate monitor processes ran for the duration of the test.

At 420,000 concurrent users, authentication latency was approximately 0.06 seconds, and message and presence latency were approximately 0.03 seconds. At 1.02 million concurrent users, authentication latency was approximately 0.05%, and message and presence latency were approximately 0.04 seconds.

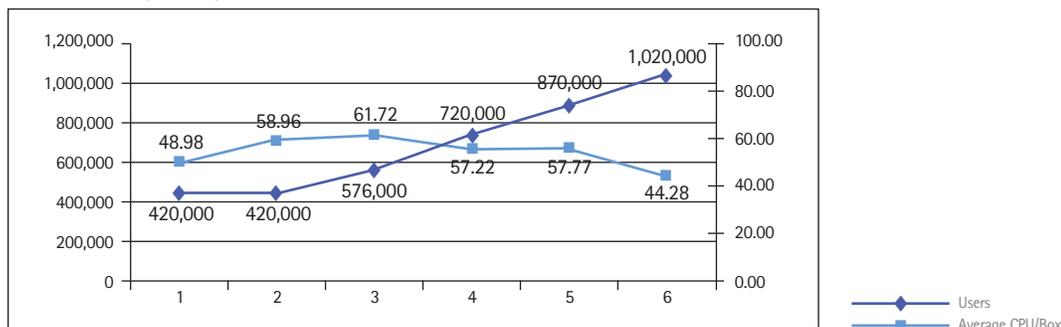


SCALABILITY ANALYSIS

Many network applications fail to scale due to their increased need to coordinate between multiple servers as additional hardware is added. Jabber XCP does not suffer from this problem, and is unique in that the load-balancing feature creates additional overhead that does not grow as the number of servers increases.

The test proved that adding a second server actually had a negative performance impact. CPU usage increased over the single server configuration due to the additional overhead created by enabling load balancing. Each additional server allowed Jabber XCP to scale by an additional 150,000 users, and the test demonstrated that CPU usage per system decreased as users were added. This indicates that more users could have been supported in this environment, further illustrating the incomparable scalability and reliability of Jabber XCP.

Scalability Analysis



CONCLUSION

The performance test results, independently verified by Sun Benchmark Centers, clearly demonstrate that Jabber XCP can handle more than a million concurrent users but, more importantly, illustrate the linear scalability of the Jabber XCP platform. The test used a traffic pattern common among providers of presence-based services, such as communications service providers, and demonstrated the highly scalable and reliable performance of Jabber XCP.

ONE MILLION USERS TEST ENVIRONMENT

To handle one million concurrent users and 1.8 million registered users, the test used 16 servers as simulators, two servers as Connection Managers, six Jabber XCP router servers (jabberd), and one database server.

USER ENVIRONMENT

- 1.02 million concurrent users
- 1.8 million registered users

SOFTWARE ENVIRONMENT

Operating System:

- Sun Solaris 10 Update 1

Database:

- Oracle 10G-R2 64-bit
- Oracle 9.2.0.7 32-bit (client on Jabber XCP machines)

Jabber XCP 5.x Simulation System:

- RedHat AS 3.0
- Jabber jLoad 1.0

HARDWARE CONFIGURATION

Jabber XCP System:

- 8 Sun T2000: 1.2GHz, 8 cores, 16 GB RAM, 4NIC cards

Database Storage System:

- Sun T2000: 1.2 GHz, 8 cores, 16 GB RAM, 4 NIC cards
- Sun StorageTek 3510 FC Array

Simulation System:

- 16 Sun AMD Dual core/hyperthreaded, 2.4 GHz

NETWORK INFRASTRUCTURE

- 2 gigabit Ethernet networks running Extreme Networks Alpine equipment
- Switches were interconnected with 10 gigabit interconnects

TEST PARAMETERS

- 30 contacts per user
- 7 active users per contact list
- 1 action per user per 120 seconds
- Memory use for all Jabber XCP components must stabilize and show no growth for a period of not less than one hour
- Average latencies for all transactions recorded by the simulator must be below two seconds

**NO OTHER
PRESENCE
PLATFORM
HAS PROVEN
THIS LEVEL OF
LINEAR SCALING**

