Sendmail, Inc.

SAMS Capacity on IBM zSeries

Sendmail Advanced Message Server

Prepared by the Sendmail Benchmarking Team and Sendmail Engineering

Edited by
Jeff Morris
Product Line Manager
April 2002

1. Table of Contents

1. TABLE OF CONTENTS 2

2. EXECUTIVE SUMMARY 3

3. SAMS POUGHKEEPSIE TEST ENVIRONMENT 4

3.1. THREE PHASES OF SAMS TESTING 4
3.2. POP VS. IMAP 4
3.3. MULTIPLE SAMS SERVERS USING PROXY 5
3.4. HOW TO READ THE RESULTS 5

4. POP ISP TEST RESULTS 6

4.1. THE POP ISP PROFILE 6
4.2. SINGLE SAMS POP, SINGLE CP 7
4.3. SINGLE SAMS POP, MULTIPLE CPS 7
4.4. MULTIPLE SAMS POP, MULTIPLE CPS 8

5. CORPORATE IMAP TEST RESULTS 9

5.1. THE CORPORATE IMAP PROFILE 9
5.2. SINGLE SAMS IMAP, SINGLE CP 10
5.3. SINGLE SAMS IMAP, MULTIPLE CPS 10
5.4. MULTIPLE SAMS IMAP, MULTIPLE CPS 11

6. CONCLUSION 12

2. Executive Summary

This document, prepared by the Sendmail Benchmarking Team, describes the results obtained in exercising the Sendmail Advanced Message Server (SAMS) on IBM zSeries Mainframes running Linux. The SAMS tests took place in IBM's Poughkeepsie, NY benchmarking center. There have been three separate testing phases, with three different goals. The third, and most ambitious phase was designed to gather enough performance and sizing information to allow IBM to make use of the LSPR (Large System Performance Reference) methodology in sizing SAMS opportunities on the Mainframe. The LSPR methodology is explained at: http://www-1.ibm.com/servers/eserver/zseries/lspr/ .

IBM and Sendmail have executed more than 200 successful capacity runs over the course of these testing exercises. The exercises follow Sendmail's Capacity Testing models described in the SAMS Capacity Testing Guide. Readers of this information should understand the methodological details described in that document.

These exercise were not designed as "Benchmarks," meaning that we were not attempting to set a performance standard. These tests were designed to capture enough information about the performance characteristics of SAMS so that we could accurately represent its capabilities to customers. We sought to answer the question, "how much mainframe do I need to run your product?" These tests provide the answer. They have also resulted in the following conclusions regarding Sendmail Advanced Message Server running on IBM's zSeries Mainframe under Linux:

- SAMS runs and scales effectively under Linux on the Mainframe in both POP and IMAP configurations.
    - Marrying SAMS multi-server capabilities with the mainframe's ability to run many Linux images across one or zSeries 900 CP, makes SAMS massively scaleable.
    - This massive scalability is achieved while maintaining a very high Quality of Service standard, by ensuring that more than 75% of transactions are completed in less than half a second and that more than 99.5% of transactions complete without error.
- SAMS POP tests indicate that a single instance of SAMS running in a single Linux image on one Mainframe CP is capable of supporting 250,000 POP users.
- SAMS IMAP tests indicate that a single instance of SAMS running in a single Linux image on one Mainframe CP is capable of supporting 6,000 simultaneously connected IMAP users.
- SAMS Multi-server POP tests indicate that three instances of SAMS, front ended with Sendmail's Message Access Proxy, running multiple Linux images across four zSeries 900 CP is capable of supporting 750,000 POP users.
    - Recognizing that current Mainframe technology has 16 available CPs, we conclude that Sendmail Advanced Message Server is capable of supporting 3,000,000 POP-ISP users on a single IBM zSeries mainframe.
    - These test results reinforce the August 2001 assertion that SAMS is capable of supporting 2,000,000 POP-ISP users on a single IBM zSeries mainframe.
- SAMS Multi-server IMAP indicate that three instances of SAMS, front ended with Sendmail's Message Access Proxy, running multiple Linux images across four zSeries 900 CP is capable of supporting 14,000 simultaneously connected IMAP users.

Each of these conclusions is explained in detail through the remainder of this document.

3. SAMS Poughkeepsie Test Environment

3.1. Three Phases of SAMS Testing

Sendmail and IBM engaged in three separate testing sessions to exercise SAMS. The objective of each test, the SAMS environment, and the mainframe used in each exercise is described in the table below.

Item	**Phase 1**	**Phase 2**	**Phase 3**
Dates of Test	May 7-25, 2001	June 18 – July 6, 2001	January 21-April 9, 2002
Objective of Test	50K POP users on 1 SAMS	400K POP users on 1 SAMS	Capacity Data for LSPR Sizing
Number of POP Runs	7	12	85
Number of IMAP Runs	14	4	92
Test Environment			
Sendmail Advanced Message Server version	1.1.1	1.1.1	1.2.2
Linux Kernel version	2.2.16	2.4.7	2.4.16
Number Linux Partitions	1	1	1-5
Linux Operating System: SuSE Version	SuSE 7.0	SuSE 7.2	SLES 7
Number of Clients: Dual P3 750mhz, RHL 7.1	4	8	8
Network: 100MB into 1GB backbone			
Workload Generator: Mozilla's MStone tool.			
Mainframe			
Model	9672-RX6	2064-116	2064-116
Processors	10	16	16
Class	G5	Freeway	Freeway
DASD: Shark F20	1	2	2

The remainder of this document outlines the findings of the Phase 3 tests, concluded in April. Conclusions here are preliminary as Sendmail and IBM study the information.

3.2. POP vs. IMAP

It is very important that the reader understand the differences in User access methods, and their effect on the mail system's performance, and their effect on the interpretation of these test results.

Sendmail usually characterizes the user load as both a function of the user's environment and their access method. Our tests are designed to gauge the capacity of SAMS in supporting ISP environments where the predominant access method is POP and in supporting corporate environments where the predominant access method is IMAP. Specific differences in characterization are outlined below, along with a summary of our testing assumptions. For specific assumptions, our rationale and our configuration of the tools that run our tests, the reader should consult the Sendmail Advanced Message Server Capacity Testing Guide.

3.3. Multiple SAMS Servers Using Proxy

We've found that Linux and other operating systems have built-in limits to them that make scaling a single SAMS server unpractical. We also seek to take advantage of the mainframe's ability to run multiple virtual Linux servers. Therefore our test exercises include many runs using multiple instances of SAMS with the SAMS Message Access Proxy directing access to the correct server.

3.4. How to Read the Results

Run Date	ID	Users Test Profile	CPs	All Pass?	SMPT 75 % cmds <200 msec	SMTP errors < .5%	POP/ IMAP 75% cmds <200 msec	75% Login times <200 msec	75% retrieval times <500 msec	POP/ IMAP errors <0.5%
Multiple SAMS POP, Multiple CPs										
4-Apr 9:35 PM	**f6**	**750K P-M**	**4S**	**Yes**	**12.4**	**0.116**	**12.2**	**161**	**164**	**0**

Each table is labeled with the category of tests run. In this table's instance, this represents data from a POP test of multiple instances of SAMS running under multiple zSeries 900 CP.

- Run Date is the date and time of the test run.
- ID is the Test Identification Number
- User Test Profile describes:
    - The number of users in the test
    - The type of user access: I=IMAP P=POP
    - The type of SAMS installation used: S=Single SAMS M=Multiple SAMS
    - Reference runs are also in bold and yellow to mark the top result for that type of test.
- CPs describes the number of zSeries 900 Computer Processors that are applied to the system.
    - D=Dedicated
    - S=Shared
- SMTP 75% cmds <200 msec is the passing Quality of Service level for SMTP commands
- SMTP errors <.5msec is the passing Quality of Service level for SMTP levels
- POP/ IMAP 75% cmds <200 msec is the passing Quality of Service level for completion of POP or IMAP commands
- 75% Login times <200 msec is the passing Quality of Service level for user login times to start a session
- 75% retrieval times <500 msec is the passing Quality of Service level for message retrieval times
- POP/ IMAP errors <0.5% is the passing Quality of Service level for POP or IMAP errors.

The following Results Tables are categorized by the Type of Test, and sorted by the number of users, and the number of CPs used to support that workload. The reference run is highlighted. Bad tests (tests that did not generate legitimate results) have been removed from these tables. Redundant tests have also been removed to improve readability of the document.

4. POP ISP Test Results

4.1. The POP ISP profile

We have assumed that ISPs who serve the consumer market prefer to provide POP user accounts.

POP users usually store mail in their mail client where their storage folders reside. Most mail is deleted from the server after it is downloaded from the user's Inbox.

POP users operate in a "hit and run" manner. POP users, or specifically ISP users, typically login, check mail, do other things, and check mail just before they log off. POP users receive all messages, regardless of size, attachment or importance, to their mail client when they log in. The affect is a time lag in which the user must wait for the completion of their mailbox transfer before they can read any message. These log in and message retrieval times are critical to understand and evaluate. We feel our acceptance criteria—that at least 75% of all logins succeed in less than 200 milliseconds, and that at least 75% of all messages retrieved and deleted from the server succeed in less than 500 milliseconds—are very strict, especially when we expect that more than 99.5% of these transactions must succeed.

When calculating POP activity, we assume that the ISP provides an access point (modem) for every ten accounts they offer. In our tests we characterize our system as having 10% of the accounts being active. Our tests are conducted assuming that these access points are saturated. Most of our current ISP customers estimate that their POP account activity rate is usually around 1-2% of their total accounts. Again, we feel that the number of active POP user accounts is high.

POP messages are typically small-ish in size, ranging from about 300 bytes to 2 megabytes, which, not coincidentally are the limits we use in our tests. The average message size in our POP tests is about 28K. That's roughly 350 lines of text. Our most popular message size is 3K, which is about a page of text, or 40 lines. Sendmail believes these are reasonable and realistic sizes. For our tests we assume that each POP account received five messages in the prior day, and will receive an average of five messages during the day in which the performance test was run. We assume that new messages arrive uniformly throughout the day and are not concentrated during specific time windows, as in our IMAP tests.

So, for our 750,000 user result documented in this paper, we have prepared and measured the following:

- We preloaded 750,000 user accounts with 5 unread messages each. That's 3,750,000 unread messages.
- When we start the test, we expect that another 156,250 messages will arrive per hour. That means that just over 20% of the accounts will receive message during the test run.
- We then spin up 75,000 simultaneously active POP user sessions, where each session lasts about 30 minutes, and each user checks mail about every 5-8 minutes during that session.
- That means that we are accommodating about 150,000 (probably a little less) user sessions per hour, and reading more than 400,000 messages per hour. Messages are read one at a time.
- Mapping back to our Quality of Service criteria, we expect that more than 112,500 of our login times succeed within two tenths of a second and more than 300,000 messages are read within half a second, and that less than .5% of these operations will fail or timeout.

Those were our assumptions. Our actual results were as follows:

- We measured 52.3 user sessions per second or 188,280 per hour. 75% of those logins averaged 161 milliseconds.

- We measured 120.53 messages read per second or 433,080 read messages per hour. 75% of these transactions happened within 164 milliseconds, and every single one was successful.
- We delivered 15.55 new messages per second, or added 55,800 new messages while the test was running. 75% of our SMTP commands succeeded in less than 13 milliseconds. We did have about 65 SMTP retry cases before their delivery succeeded, accounting for our .116% SMTP command error rate.
- Our average CPU utilization rate across our four zSeries 900 CP was 85.66% during the peak period. Network, IO and storage bandwidth were ample for the test.

While we continue to assert that these tests were not benchmarks, we are confident in our claims that a properly configured mainframe could handle at least 3 million users on the system, provided that these users have ample network bandwidth and I/O.

Sendmail continues to investigate the optimal combination of hardware, networking and SAMS configuration in which to attempt an Industry record SpecMail benchmark. In doing so, we seek to not only break the performance records established by other vendors, but to do so on a hardware topology that costs a fraction of what these vendors have used.

4.2. Single SAMS POP, Single CP

Run Date	ID	Users Test Profile	CPs	All Pass ?	SMPT 75 % cmds <200 msec	SMTP errors < .5%	POP/ IMAP 75% cmds <200 msec	75% Login times <200 msec	75% retrieval times <500 msec	POP/ IMAP errors <0.5%
Single SAMS POP, Single CPs										
15-Feb 2:11 AM	26	100K P-S	1D	Yes	3.59	0	0.48	47.2	38.8	0
19-Feb 11:41AM	33	200K P-S	1D	Yes	7.6	0	0.7	61	99.3	0
12-Mar 12:02AM	94	225K P-S	1S	Yes	9.69	0	0.71	44.5	70	0
19-Feb 3:57 PM	34	250K P-S	1D	Yes	12.6	0.002	0.93	88	224	0
19-Feb 8:21 PM	**35**	**250K P-S**	**1D**	**Yes**	**13**	**0.022**	**1.03**	**87.1**	**218**	**0**
20-Feb 1:36 AM	36	275K P-S	1D	No	22.1	0.086	3.56	130	1,420	0.05

4.3. Single SAMS POP, Multiple CPs

Run Date	ID	Users Test Profile	CPs	All Pass ?	SMPT 75 % cmds <200 msec	SMTP errors < .5%	POP/ IMAP 75% cmds <200 msec	75% Login times <200 msec	75% retrieval times <500 msec	POP/ IMAP errors <0.5%
Single SAMS POP, Multiple CPs										
20-Feb 5:21 PM	37	250K P-S	2D	Yes	3.48	0	0.7	64.6	107	0
26-Feb 1:29 AM	48	250K P-S	2D	Yes	3.98	0	0.92	80.6	142	0
26-Feb 4:11 PM	**49**	**250K P-S**	**2D**	**Yes**	**3.58**	**0**	**0.78**	**60.7**	**101**	**0**
26-Feb 10:24PM	50	300K P-S	2D	No	4.95	0	1.93	150	6Sec	0.08
21-Feb 1:39 AM	38	350K P-S	2D	No	0.73	0	4.74	139	14Sec	0.03

4.4. Multiple SAMS POP, Multiple CPs

Run Date	ID	Users Test Profile	CPs	All Pass ?	SMPT 75 % cmds <200 msec	SMTP errors < .5%	POP/ IMAP 75% cmds <200 msec	75% Login times <200 msec	75% retrieval times <500 msec	POP/ IMAP errors <0.5%
Multiple SAMS POP, Multiple CPs										
28-Feb 3:01 PM	56	300K P-M	2S	Yes	3.01	0	1.4	47.3	73.7	0.06
28-Feb 6:28 PM	57	350K P-M	2S	Yes	3.12	0	1.73	61	119	0.05
28-Feb 9:18 PM	58	375K P-M	2S	Yes	3.25	0	2.02	72.9	194	0
9-Apr 9:30 AM	**g6**	**400K P-M**	**2S**	**Yes**	**13.4**	**0**	**8.11**	**87.6**	**99.2**	**0**
1-Mar 5:03 AM	60	450K P-M	2S	No	4.42	0	10.7	153	2Sec.	0.04
3-Apr 2:38 PM	**f2**	**600K P-M**	**3S**	**Yes**	**20.1**	**0.192**	**13.1**	**133**	**149**	**0**
15-Mar 8:02 AM	a1	675K P-M	3S	No	8.52	0	5Sec	13Sec	14Sec	0
5-Mar 12:36 AM	78	450K P-M	4D	Yes	4.37	0	1.33	97.3	46.2	0.03
5-Mar 3:01 PM	80	600K P-M	4D	Yes	5.49	0	2.84	67.4	82.4	0
13-Mar 5:23 AM	96	600K P-M	4S	Yes	3.19	0	1.88	55.3	81.8	0
13-Mar 2:13 PM	97	675K P-M	4S	Yes	3.48	0	5.37	101	132	0
4-Apr 9:35 PM	**f6**	**750K P-M**	**4S**	**Yes**	**12.4**	**0.116**	**12.2**	**161**	**164**	**0**
14-Mar 2:17 AM	99	825K P-M	4S	No	6.58	0.011	5Sec	13Sec	20Sec	0.81

5. Corporate IMAP Test Results

5.1. The Corporate IMAP Profile

IMAP is a the preferred protocol for Corporate users because it allows IT to centralize many important operations like backups and quotas while still providing each user optimal flexibility with regard to how they get their mail.

IMAP users usually store all their mail on the server, where their IMAP folders reside. Mail is usually not deleted. IMAP users have the benefit of accessing their mailbox from a number of different mail clients including traditional desktop clients, webmail clients, PDAs and wireless phones. IMAP's other big advantage is the speed at which messages are delivered to the client. IMAP is faster because it only retrieves parts of each message after the user completes authentication. The result is the user interacts with their mailbox almost immediately. The completion of the message retrieval happens when the user chooses to read the message, or to view an attachment. In our tests, the IMAP commands measurement accounts for the delivery of email headers, while the retrieval time accounts for the delivery of each message body.

We assume that corporate IMAP users receive more mail, and their mail is usually larger in size than POP users. Also, IMAP users login and stay connected for a much longer period of time—often all day. So, for our tests, we assume that every IMAP user is connected to SAMS. Our intent is to simulate the highest activity level during a day—the morning email check. We know that in real life not all users are always connected, but in order to develop realistic sizing we always test with 100% concurrency. For our tests, IMAP users received 50 messages in their inbox. IMAP users' message size is somewhat larger, following the same distribution method as the POP users—300bytes are the smallest messages, 2MB are the largest. 4K, or just over 50 lines, is our most popular message size, while 37K is our average—about 10 pages of text.

We anticipate that 95% of new messages arrive during business hours. So our new message arrival rate is similar to the POP tests. Examining these assumptions against our Single IMAP server running within a single zSeries 900 CP, we assume:

6,000 user accounts with 50 messages in their inbox. So we preload 300,000 messages. We assume that these users will receive 90% of their mail throughout the 8-hour workday. So we expect the system to generate about 33,750 messages per hour. We also expect all users to log in during the 8am hour and to stay connected for the duration of the test. Their check interval is assumed to be every 5 to 8 minutes.

Our Quality of Service standards are the same as for POP. 75% of logins succeed within 200 milliseconds while the same percentage of message retrievals succeed within a half second. Remember our average message size is 33% larger than that of the POP user.

Our results are as follows:

- We measured 1.67 user sessions per second or all 6,000 users became active within the first hour. 75% of those logins averaged 64.7 milliseconds. (.065 seconds!)
- We measured 18.44 messages read per second or 66,384 read messages per hour. 75% of these transactions happened within 297 milliseconds. Our IMAP commands succeeded 99.85% of the time.
- We delivered 3.47 new messages per second, or added 12,492 new messages while the test was running. 75% of our SMTP commands succeeded in less than 11 milliseconds, with no errors.

Here, we have proven the durability of our IMAP server under heavy load.

5.2. Single SAMS IMAP, Single CP

Run Date	ID	Users Test Profile	CPs	All Pass ?	SMPT 75 % cmds <200 msec	SMTP errors < .5%	POP/ IMAP 75% cmds <200 msec	75% Login times <200 msec	75% retrieval times <500 msec	POP/ IMAP errors <0.5%
Single SAMS IMAP, Single CPs										
15-Feb 2:28 PM	27	3.5K I-S	1D	Yes	7.13	0	20.6	27.1	91.2	0
15-Feb 7:44 PM	29	4K I-S	1D	Yes	7.63	0	25.1	32.8	112	0.3
2-Mar 11:28 PM	71	5K I-S	1D	Yes	10.3	0	51.1	42.6	218	0.02
3-Mar 7:07 PM	**73**	**6K I-S**	**1D**	**Yes**	**10.7**	**0**	**81.2**	**67.9**	**297**	**0.15**
4-Mar 11:54 AM	75	7K I-S	1D	No	14.9	0	164	133	588	0.14

5.3. Single SAMS IMAP, Multiple CPs

Run Date	ID	Users Test Profile	CPs	All Pass ?	SMPT 75 % cmds <200 msec	SMTP errors < .5%	POP/ IMAP 75% cmds <200 msec	75% Login times <200 msec	75% retrieval times <500 msec	POP/ IMAP errors <0.5%
Single SAMS IMAP, Multiple CPs										
1-Mar 9:45 PM	64	8K I-S	2D	Yes	6.18	0	45.3	47.8	211	0.05
2-Mar 4:44 AM	67	9K I-S	2D	Yes	6.86	0	60.8	63	274	0.06
2-Mar 8:19 AM	**69**	**10K I-S**	**2D**	**Yes**	**8.25**	**0**	**86.2**	**85.9**	**420**	**0.07**
9-Feb 6:23 PM	14	10K I-S	3D	Yes	3.9	0	49.2	54.1	239	0.04
1-Feb 1:12 AM	1	7K I-S	2S	Yes	5.79	0	35.7	42.2	143	0.02
4-Feb 7:32 PM	5	10K I-S	3S	Yes	2.62	0	53.4	55.3	229	0.05
30-Jan 12:32 AM		3K I-S	4S	Yes	2.12	0	9.11	16.3	48.2	0.23
30-Jan 9:48 PM	Y	9K I-S	4S	Yes	2.54	0	42.3	43.4	191	0.022
31-Jan 3:13 AM	Z	14K I-S	4S	No	2.86	0.003	157	82.9	1,730	0.417

5.4. Multiple SAMS IMAP, Multiple CPs

Run Date	ID	Users Test Profile	CPs	All Pass ?	SMPT 75 % cmds <200 msec	SMTP errors < .5%	POP/ IMAP 75% cmds <200 msec	75% Login times <200 msec	75% retrieval times <500 msec	POP/ IMAP errors <0.5%
Multiple SAMS IMAP, Multiple CPs										
6-Feb 9:43 PM	9	5K I-M	2S	Yes	18	0	89.9	125	347	0.48
1-Mar 3:03 PM	61	6K I-M	2S	Yes	3.02	0	16.9	33.8	74.5	0.78
1-Mar 6:56 PM	63	8K I-M	2S	Yes	3.99	0	30	48.7	123	0.07
8-Apr 4:04 PM	g4	8K I-M	2S	OK	15.5	0	144	193	515	0.37
1-Mar 11:16 PM	65	8.5K I-M	2S	Yes	4.8	0	39.1	55.5	154	0.35
2-Mar 2:39 AM	66	9K I-M	2S	Yes	4.95	0	39	55.4	150	0.25
2-Mar 6:43 AM	68	10K I-M	2S	Yes	7.86	0	56.6	70.9	212	0.24
2-Apr 8:43 PM	f0	9K I-M	3S	Yes	12.5	0	82.9	135	316	0.13
3-Apr 1:52 AM	**f1**	**10.5K I-M**	**3S**	**OK**	**12.3**	**0**	**165**	**171**	**512**	**0.15**
17-Mar 4:03 AM	b3	12K I-M	3S	Yes	23.3	0.01	52.8	75.4	183	0.29
18-Mar 1:15 AM	b5	13.5K I-M	3S	Yes	83.8	0.136	103	102	378	0.09
18-Mar 5:49 AM	b7	14.4K I-M	3S	Yes	80.8	0.135	91.5	119	314	0.16
7-Feb 6:24 PM	10	7K I-M	4D	Yes	11	0	79.1	93.8	321	0.01
8-Feb 1:53 AM	11	8K I-M	4D	Yes	10.8	0	106	133	489	0.07
6-Apr 7:08 PM	g1	13.5KI-M	4S	Yes	9.87	0	92.4	136	333	0.1
7-Apr 5:53 PM	**g3**	**14.4K I-M**	**4S**	**OK**	**10.6**	**0**	**156**	**204**	**549**	**0.31**
12-Feb 4:21 PM	18	8K I-M	6D	Yes	9.68	0	60.5	73.5	260	0.25
12-Feb 6:18 PM	19	9K I-M	6D	Yes	10.4	0	75.1	88	323	0.38
13-Feb 2:06 AM	21	10K I-M	6D	Yes	10.7	0	93.5	105	391	0.24
13-Feb 11:26AM	22	11K I-M	6D	Yes	10.6	0	107	129	484	0.18
13-Feb 1:14 PM	23	12K I-M	6D	No	10.7	0	123	139	599	0.23

6. Conclusion

The results published here demonstrate the scalability of Sendmail Advanced Message Server. As a POP access server, it has been demonstrated that SAMS can handle 750,000 users on only four zSeries 900 CP. Simple arithmetic tells us that this is capable of supporting 3,000,000 heavy POP users, provided that the system has ample storage and network bandwidth to accommodate the load.

The reader should notice that none of these test results describe any server failures. And they illustrate the strict Quality of Service Standard that Sendmail applies in evaluating the success or failure of the SAMS test.

SAMS on the mainframe is massively scaleable, and utterly reliable. For large scale, mission-critical email, there is no better alternative than Sendmail on zSeries.