AMS Note AMSnote-2003-10-01

Oct. 10, 2003

Data Transmission Test for AMS-02 SOC at Southeast University

Jian Gong, Hua Wu, Guang Chen, Jianzhang Wu

Abstract

A data transmission test has been conducted between SEU and CERN for the AMS-02 SOC establishment at SEU, Nanjing. The test including its environment is introduced and some test results are analyzed in this report.

1 Introduction

SEU has joined the AMS-02, an international cooperative project, since the middle of 2002. According to the arrangement, A SOC (regional center at SEU, RC@SEU for short) will be established in SEU before Oct., 2004. As a result of the meeting held in the early Nov., 2002 between SEU and CERN in Nanjing, some link tests and data transmission tests were carried out between Nanjing and Geneva to find out if there are enough bandwidth available to meet the need of RC@SEU. The tests were conducted during the last week of November and the first week of December, 2002 at both sites separately. The test environment and the test conducted by SEU side are presented in this report.

2 The test environment

This test is conducted in China Education and Research Network (CERNET) Eastern China (North) Regional Network Center, which located in the campus of SEU, and the test network is attached to the CERNET regional backbone directly. As required by AUP of CERNET, a special IP block was applied from CERNET network center for this test. The Final RC@SEU network will be moved into the SEU campus network, which has a GE based switching infrastructure, so a few more hops added will not effect the throughput significantly.

The test machine, amsseu.njnet.edu.cn is configured with

- Linux version 2.4.18-14(Red Hat Linux 8.0 3.2-7),
- 256M RAM,
- CPU: Intel(R) Pentium(R) 4 CPU 1.80GHz,
- 60G disk, and
- 100M network controller.

BBFTP2.2.1 and some other tools required by the test are installed. Another similar configured PC is used in the test environment for the local testing purpose.

3 Link analysis

We used *traceroute* to look at the path between amsseu.njnet.edu.cn and pcamsf0.cern.ch, and found that the route between two sites is asymmetric. The path from Nanjing to Geneva consists of 13 hops, but the reverse path consists of 25 hops (see table 1 and 2). It seems that the Startap treat these packets from Geneva to Nanjing as ordinary traffic and route them to the normal Internet, instead of sending them the CERNET link directly. The test shows that the delay is roughly the same, this is because *traceroute* is not the suitable tool to test the one-way delay, some other test method should be used if delays for each direction are wanted to be tested accurately. The longer the path, the easier the performance could be effected by other unrelated traffics.

We discussed the situation with CERNET network center after the test, and the routing configuration had been modified at Startap to make the route symmetric.

Table 1 Traceroute results from amsseu.njnet.edu.cn to pcamsf0.cern.ch

	Hop	Max	Ave.	Min	Comment
1	210.25.240.1 (210.25.240.1)	0.388 ms	0.316 ms	0.305 ms	Nanjing
2	202.112.24.21 (202.112.24.21)	0.177 ms	0.096 ms	0.099 ms	
3	nj0.cernet.net (202.112.53.85)	0.160 ms	0.136 ms	0.137 ms	
4	whnj4.cernet.net (202.112.46.73)	8.684 ms	8.642 ms	8.598 ms	Wuhan
5	bjwh4.cernet.net (202.112.46.65)	23.850 ms	23.807 ms	23.824 ms	Beijing
6	202.112.61.194 (202.112.61.194)	23.956 ms	23.924 ms	23.911 ms	
7	202.112.61.22 (202.112.61.22)	207.300 ms	207.252 ms	207.449 ms	Stockten
8	pass.bjnet.edu.cn (202.112.61.6)	306.037 ms	287.698 ms	269.743 ms	Chicago
9	m10-m5.startap.net (206.220.240.106)	276.155 ms	287.381 ms	262.957 ms	
10	ar3-chicag0-startap.cern.ch (192.65.184.246)	279.480 ms	279.652 ms	279.485 ms	
11	cernh4-stm4.cern.ch (192.65.184.26)	412.145 ms	410.637 ms	410.659 ms	Geneva
12	cernh2.cern.ch (192.65.185.2)	423.172 ms	411.032 ms	411.432 ms	
13	pcamsf0.cern.ch (137.138.206.126)	412.258 ms	411.899 ms	412.262 ms	

Table 2 Traceroute results from pcamsf0.cern.ch to asmseu.njnet.edu.cn

	Hop	Max	Ave.	Min	Comment
1	r892-s-rca80-1-ip12 (137.138.206.65)	0.208 ms	0.165 ms	0.142 ms	Geneva
2	b887-r-rca86-1-gb18 (194.12.129.45)	0.404 ms	0.253 ms	0.232 ms	
3	b513-b-rca86-2-rb7 (194.12.132.25)	0.681 ms	0.475 ms	0.317 ms	
4	b513-e-rca80-1-pg2 (194.12.133.250)	0.636 ms	0.583 ms	0.313 ms	
5	cernh2-dmz (192.65.184.65)	1.369 ms	1.015 ms	0.792 ms	
6	cernh4 (192.65.185.4)	1.062 ms	0.793 ms	0.772 ms	
7	ar3-chicag0-stm4 (192.65.184.25)	132.127 ms	131.973 ms	131.933 ms	Chicago
8	ar1-chicago-fe (192.65.184.116)	133.463 ms	131.951 ms	131.947 ms	
9	NYC-gw14.NYC.US.net.DTAG.DE (62.156.138.190)	122.142 ms	122.426 ms	121.911 ms	New York
10	NYC-gw15.USA.net.DTAG.DE (62.156.131.158)	128.198 ms	122.244 ms	123.010 ms	
11	sl-gw31-nyc-12-0.sprintlink .net (144.223.27.133)	122.244 ms	122.157 ms	122.093 ms	
12	sl-bb27-nyc-12-0.sprintlink .net (144.232.13.35)	122.578 ms	122.231 ms	122.246 ms	
13	sl-bb26-nyc-8-0.sprintlink. net (144.232.7.37)	122.834 ms	122.289 ms	122.346 ms	
14	sl-bb23-pen-12-0.sprintlink .net (144.232.20.95)	180.707 ms	126.614 ms	126.861 ms	Pennsylv ania??

15	sl-bb22-pen-14-0.sprintlink .net (144.232.8.178)	128.559 ms	128.183 ms	128.473 ms	
16	sl-bb21-stk-10-3.sprintlink .net (144.232.18.94)	188.535 ms	187.269 ms	187.235 ms	Stockten
17	sl-gw26-stk-9-0.sprintlink. net (144.232.4.105)	187.399 ms	186.912 ms	187.088 ms	
18	sl-crnt-2-0.sprintlink.net (144.223.71.222)	207.131 ms	206.370 ms	206.453 ms	??
19	202.112.61.21 (202.112.61.21)	389.153 ms	388.779 ms	388.320 ms	Beijing
20	202.112.61.193 (202.112.61.193)	389.919 ms	388.508 ms	388.645 ms	
21	whbj4.cernet.net (202.112.46.66)	404.067 ms	404.148 ms	404.562 ms	Wuhan
22	njwh4.cernet.net (202.112.46.74)	412.169 ms	411.966 ms	412.060 ms	Nanjing
23	nj1.cernet.net (202.112.53.86)	413.171 ms	412.116 ms	412.701 ms	
24	202.112.24.22 (202.112.24.22)	412.706 ms	413.222 ms	412.532 ms	
25	210.25.240.10 (210.25.240.10)	412.710 ms	412.102 ms	412.360 ms	

4 Bandwidth tests

4.1 Tests using Iperf

Iperf is a tool to measure maximum TCP bandwidth along a path, allowing the tuning of various parameters and UDP characteristics. On Dec.3, 2002, Iperf was used to measure the TCP bandwidth available between pcamsf0.cern.ch (as client) and amsseu.njnet.edu.cn (as server). The test showed that the bandwidth grew as the number of parallel TCP streams increased up to 2.9Mb/s. Because the routes between two sites are asymmetric, the bandwidths for two sites are slightly different. Table 3 shows that the more the numbers of TCP connection, the better the bandwidth available from the client side than the server side. This implies that there might be a bottleneck existed in Startap because from Geneva to Chicago, and from Stockten to Nanjing are shared links, from Stockten to Chicago being only one hop.

Table 3 Bandwidth tested with Iperf

# of TCP connection	Bandwidth for Server (Kb/s)	Max Bandwidth for Server (Kb/s)	Bandwidth for Client (Kb/s)	Max Bandwidth for Client (Kb/s)
1	923	955	812	884
2	1889	1977	1851	1984
3	2840	2868	2834	2914
4	2211	2423	2244	2473
5	2828	2948	2933	2948
6	2056	3012	2254	3012

4.2 Test results of Pathchar

Pathchar is a testing tool, written by Van Jacobson at Lawrence Berkeley Laboratory (LBL), that

tries to infer the characteristics of individual links along an Internet path by measuring the round trip time, throughput and some other features of packets sent from a single host. With *pathchar*, four tests had been done, dedicating to find the bandwidth available along the path between SEU and CERN. The results were summarized in table 4.

Table 4 Bandwidth test with Pathchar

Hop	name	1st test	2nd test	3rd test	4th test	Comment
0	localhost	55 Mb/s	57 Mb/s	55 Mb/s	53 Mb/s	
1	210.25.240.1 (210.25.240.1)	n/a	n/a	n/a	n/a	
2	202.112.24.21 (202.112.24.21)	396 Mb/s	825 Mb/s	327 Mb/s	n/a	
3	nj0.cernet.net (202.112.53.85)	2475 Mb/s	853 Mb/s	512 Mb/s	n/a	
4	whnj4.cernet.net (202.112.46.73)	1664 Mb/s	8960 Mb/s	n/a	n/a	
5	bjwh4.cernet.net (202.112.46.65)	157 Mb/s	896 Mb/s	256 Mb/s	171 Mb/s	
6	202.112.61.194 (202.112.61.194)	7936 Mb/s	52 Mb/s	37 Mb/s	24 Mb/s	
7	202.112.61.22 (202.112.61.22)	262 Kb/s	600 Kb/s	n/a	n/a	Bottleneck?
8	pass.bjnet.edu.cn (202.112.61.6)	295 Kb/s	17.0 Kb/s	n/a	n/a	Bottleneck?
9	m10-m5.startap.net (206.220.240.106)	56.8 Kb/s	n/a	n/a	n/a	Bottleneck?
10	ar3-chicag0-startap.cern.ch (192.65.184.246)	14 Mb/s	n/a	n/a	n/a	
11	cernh4-stm4.cern.ch (192.65.184.26)	19 Mb/s	7.6 Mb/s	n/a	n/a	
12	cernh2.cern.ch (192.65.185.2)	14 Mb/s	16 Mb/s	n/a	n/a	
13	pcamsf0.cern.ch (137.138.206.126)	23 Mb/s	23 Mb/s	26 Mb/s	39 Mb/s	

Once again, the above results confirm us that the bandwidth bottleneck probably exists in Startap, especially around the connections with CERNET (a Startap managed unknown type router).

4.3 Test results of BBFTP

The bbftp tests were carried out between amsseu.njnet.edu.cn and pcamsf0.cern.ch by both SEU and CERN sides. At SEU side, the test was done locally at first to allow us to get familiar with the usage of BBFTP. The results are shown in table 5.

Table 5 Test of BBFTP in local environment

NumberOfParallelStreams	compression(gzip)	SSL	file size(Mbytes)	transmit rate（Mbits/s)
1	no	no	255.8	86.2/87.7/87.4
3	no	no	255.8	63.0/67.2/63.3
5	no	no	255.8	64.8/63.6/67.2
1	yes	no	255.8	35.5/35.1/35.0
3	yes	no	255.8	38.3/37.5/37.2
5	yes	no	255.8	39.5/39.9/39.0

1	no	yes	255.8	87.8/87.3/87.7
3	no	yes	255.8	63.2/65.0/67.0
5	no	yes	255.8	68.1/67.5/66.6
1	yes	yes	255.8	32.2/36.9/35.0
3	Yes	yes	255.8	38.3/38.2/38.0
5	Yes	yes	255.8	38.4/39.8/40.8
1	No	yes	2097152	87.1 bbftp cost 184s ftp cost 250s

The local test showed that the performance of BBFTP was much better than the ordinary FTP we used before, e.g. Wu-ftp; and the data compression and encryption would reduce the performance greatly.

Two sample files with the length of 14MB and 400MB each were used to conduct the real BBFTP tests between SEU and CERN, using 15 TCP streams. The summarized results are given in table 6, and the detailed results are shown in figure 1.

Table 6 Test between CERN and Southeast University

File size (Bytes)	Average Transfer rate (Mbits/s)	MAX Transfer Rate (Mbits/s)	MIN Transfer rate (Mbits/s))
14172160	4.28	5.01	3.66
419430400	14.575	15.9	12.7

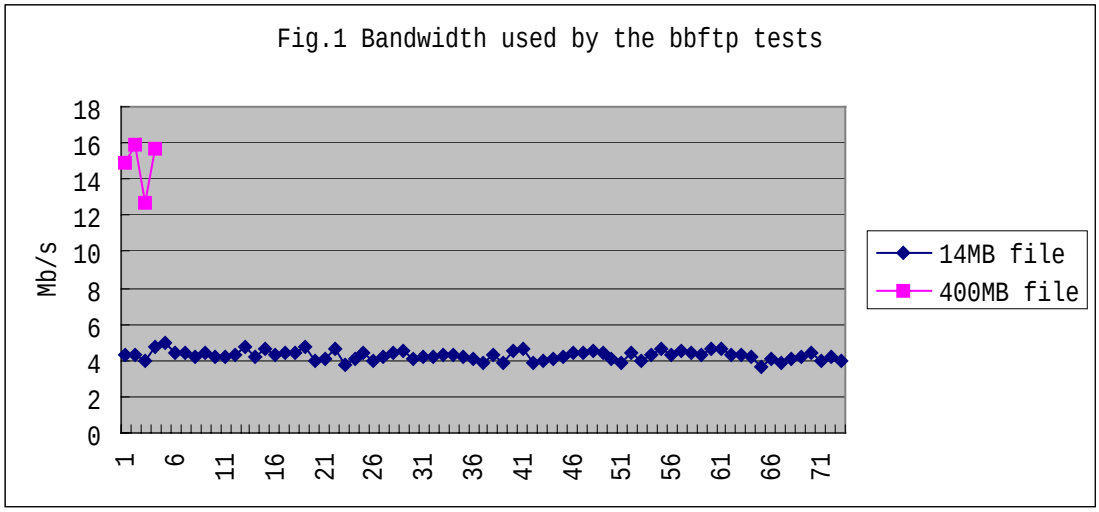

Compared with the similar tests already done by CERN, and the foreseen bandwidth requirement for the AMS-02 [2][3], the Internet links between SEU and CERN should have sufficient bandwidth to meet the needs of RC@SEU for data transmission and some other networking services.

5 Conclusions

Concerning this link testing, three things are worth to note. Firstly, from the previous tests we can

find that the path between SEU and CERN was asymmetric, and it should be corrected to make the path more stable, although currently there is no evidence that this asymmetry may effect the delay and throughput of the path between two sites.

Secondly, several tools were used to test the link performance between SEU and CERN, and different results were achieved. The tests showed that a bandwidth bottleneck might exist around the connection between CERNET and Startap, so that more testing and measurement are suggested to explore the problem further.

According to the bbftp testing, the average bandwidth available for data transmission is 4.28Mb/s for small files, and 14.575Mb/s for large files, with 15.9Mb/s as the maximum bandwidth available for large files, so that the networking path between SEU and CERN could meet the needs of data transmission for RC@SEU in principle.

Acknowledgement

This test was carried out with the help of Dr. Alexei Klimentov of AMS-02 Computing Group at CERN. He provided us the BBFTP software, and also the test environment at CERN side. Professor Xing Li of Tsinghua University at CERNET network center helped us to establish the test path between SEU and CERN with the special allocated IP address block.

References

[1] http://www.startap.net/ENGINEERING/TECHINFO.html.

[2] P.Fisher, A.Klimentov, A.Mujunen, J.Ritakari, AMS Ground Support Computers For ISS mission, AMS Note 2002-03-01, March 12, 2002.

[3] A.Elin, A.Klimentov, Data transmission program for the AMS-02 ISS Mission, AMS Note 2001-11-02, Aug.14,2002.