TS-1 Performance and Load Test Results

Introduction

This document describes a series of tests that were conducted on a TS-1 machine. The goal of these tests was to measure the limits of a TS-1 unit.

TS-1 is a fast, easy and cost–effective solution for a business PBX. It is a solid state Linux machine preloaded with Asterisk. The device contains the following components:

- Rapid Xorcom 1.1.1
- Asterisk 1.0.10
- Linux 2.6.12
- Asterisk Management Portal 1.0.10
- CPU: VIA Nehemiah 1000.735 MHz, 64KB cache
- Memory: 248632KB (8MB allocated to the built-in video card)

Worst Case Scenario

This first test was conducted to determine the amount of simultaneous calls that the TS-1 can handle at the same time. We measured the maximum number of calls generated using 4 different codecs: G.711, GSM, G.729 and Speex. We used SIP as the transfer protocol.

The TS-1 played a 60-minute WAV file (encoded as 8khz 16 bit mono), into a dedicated SIPP server, which acted as an echo server, in order to emulate a real call. The audio was read as linear (the input from the WAV file simulated the SIP phone) and sent the call to another party (the SIPP server acted as another SIP phone). Every 15 seconds another new call was added until the TS-1 stopped responding to keyboard input.

TS-1 generates calls to a SIPp server, as many calls as it can handle, and the SIPp server echoes them back to the TS-1.

We used the load average to measure “how much the CPU was stressed”. For the “best case” scenarios, a good load average is “1”. In our tests, we found that machines can behave quite well, even under load averages of “5”. At the end of the tests we saw a big jump in load average, which we interpreted as “the server cannot handle new calls”.

Worst Case Scenario – Results

ULAW/ALAW Encoding. The TS-1 was capable of handling 107 calls before the load average jumped to 24. For call 108 the load average was 51, and for call 109 (during which the test was halted) the load average was 82.

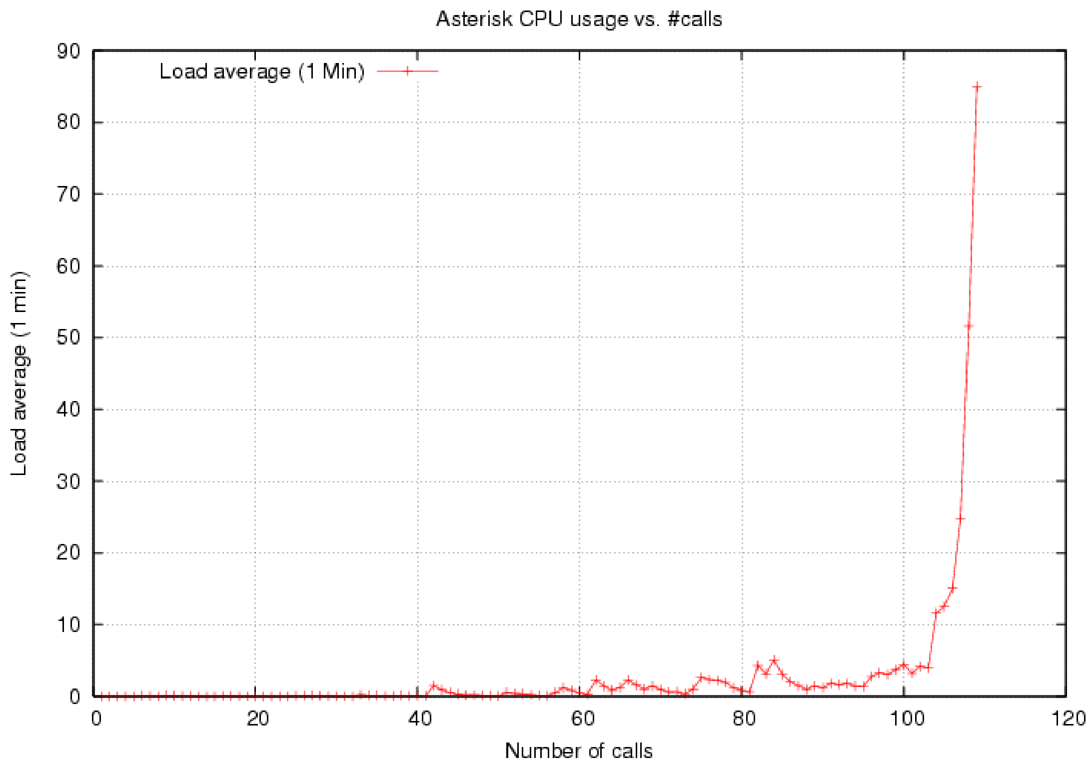

Conclusion: The TS-1 can handle 103 concurrent G.711 calls.

GSM Encoding. In this test, the load average measured for 23 calls (when the test was halted) was 20.

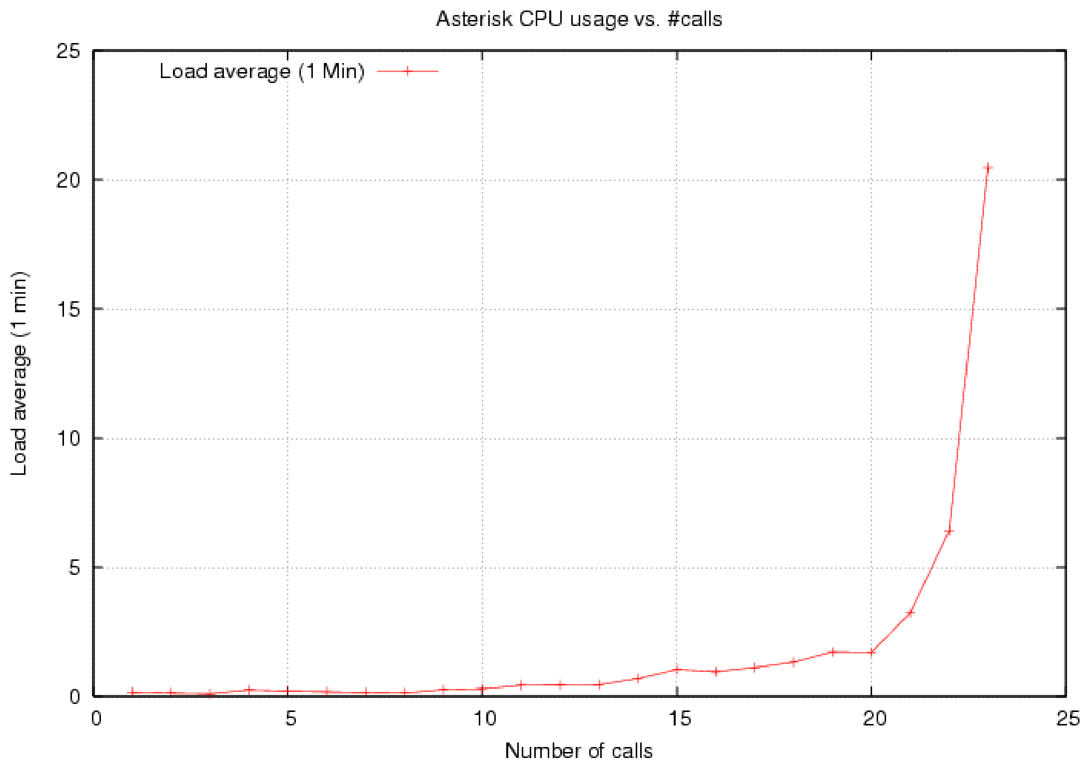

Conclusion: The TS-1 can handle 21 concurrent GSM calls.

G729 Encoding. For this test we used the codecs supplied from Digium and not the open source alternatives. For 15 calls we found that the load average was 17.

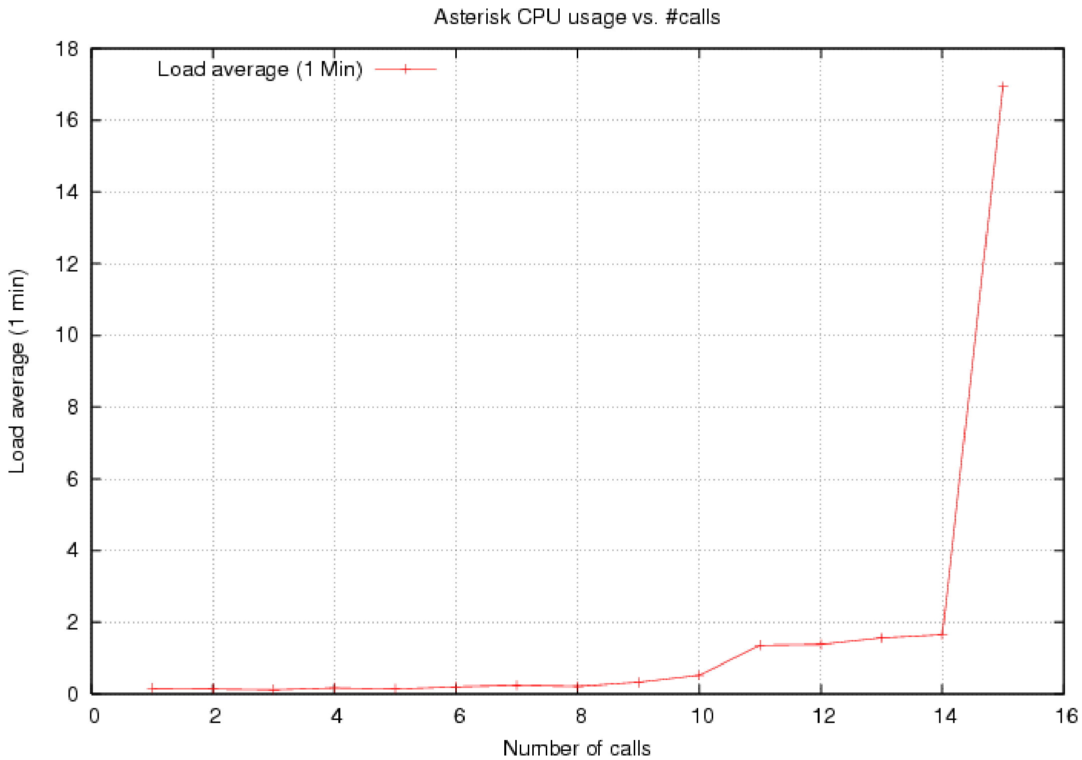

Conclusion: The TS-1 can handle 14 concurrent G.729 calls.

Speex Encoding. For the last test we measured the free codec available with the default Asterisk installation. This codec provides the best audio compression and the least bandwidth consumption, but is the one which requires the most CPU for processing. We measured a load average of 25 on 8 calls.

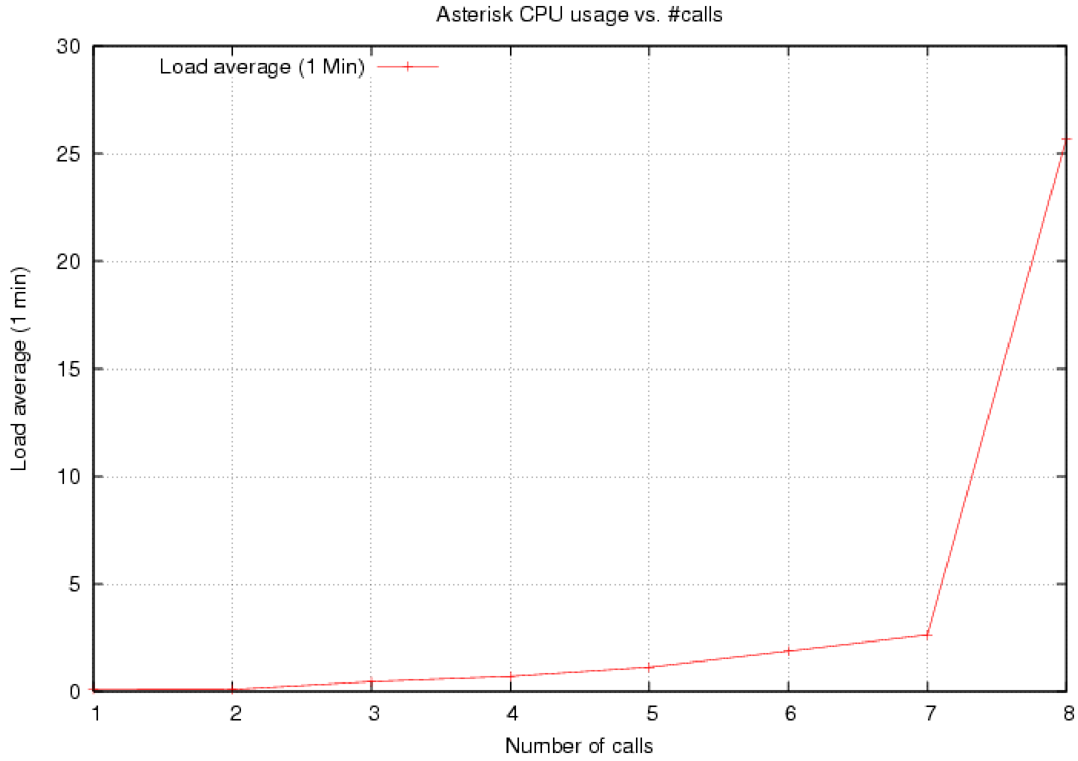

Conclusion: The TS-1 can handle 7 concurrent Speex calls.

Worst Case Scenario – Conclusion

From the data collected, we found that the TS-1 can transcode up to 103 calls using G.711, 20 GSM calls, 14 G.729 calls and 7 Speex calls.

In real life, the TS-1 can handle many more calls when some of the calls are handled inside the LAN as G.711 calls. The Asterisk issues a peer-to-peer call from one IP Phone to another – so the Asterisk is routing the voice traffic.

We saw that compiling the kernel for our architecture gives us 5 more G.711 calls than a generic i386 kernel provided by Debian. We could not measure any change on other codecs. The kernel used in this test was compiled for our architecture.

At the beginning of the tests we used Asterisk 1.0.9. With the older version of Asterisk we found that the maximum capacity of the TS-1 was 93 G.711 calls. With the new version the number of calls went up by more than 10 new calls.

Steady State Test

In this test we performed a lengthy stress test on the TS-1. The test configuration was based on another Asterisk server generating hour-long calls to the tested TS-1, and the tested TS-1 routing the SIP call to another Asterisk, serving as a gateway between 2 SIP networks. This object of the test was to check the long term stability of the system and memory leakage, and to check whether the load average rises over time.

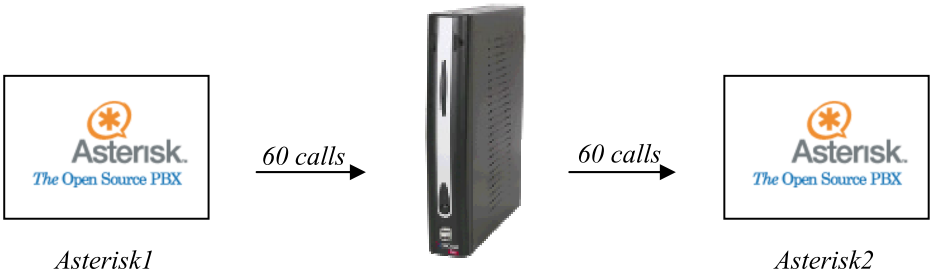

The Asterisk on the left generates a call into the TS-1, and the TS-1 forwards the call to the Asterisk on the right. The TS-1 serves as a gateway between 2 SIP networks.

The test was generated automatically via *cron* tasks and the results were emailed every hour. Another email was issued on a daily basis to summarize the day. The codec in use was G.711. The full scenario was:

1. *Asterisk1* (on the left) generates 6 calls into the TS-1.
2. The TS-1 calls *Asterisk2* (on the right) and routes the incoming call from *Asterisk1* to *Asterisk2*.
3. TS-1 measures the load average, memory consumption and number of file descriptors by the Asterisk process and sends this via email to a dedicated server.
4. After 10 minutes *Asterisk1* generates 6 new calls to the TS-1.
5. After an hour of work, every 10 minutes 6 calls are terminated and 6 new calls are generated, which gives us a steady state of 60 incoming calls and 60 outgoing calls in the TS-1 gateway.

In the original tests we used a WAV file which was one hour long, but in the final test (for which the results are presented here) we used a 10-minute file, which was looped with a dial-plan configuration 6 times. This was done in order to conserve memory, allowing a smaller file to be loaded into the memory. We also tested this scenario with 5-second files, but found that Asterisk was spending too much CPU doing the software loop.

The test described in this document was performed during the entire month of January 2006, and demonstrated no memory leaks.