Summary Test Report
Acme Packet Net-Net PAC
Performance, Availability & Capacity

Date of report: July15, 2004
Test engineer: Jason Sargeant
Report revision: 1.0

Dates of tests: June 15 – 17, 2004
Report Authors: Jason Sargeant / Carla Wandell

1 Test Overview

CT Labs was contracted by Acme Packet™ to verify the performance of the Acme Packet Net-Net PAC™ session border controller and to produce a report summarizing the results. To facilitate this Empirix provided testing products, script engineering, and lab facilities to support this high call volume test. During the test, quality assurance engineers from Acme Packet were available to help configure and support their product.

2 Product Description

The recently unveiled Acme Packet Net-Net PAC is a scalable, full-featured session border control solution designed to satisfy the performance, availability and capacity requirements of large-scale Tier 1 service provider deployments. It is comprised of a pack of up to nine 1U Net-Net session border controllers – one Net-Net Session Router (SR) and 8 Net-Net Session Directors (SD) - that performs as one logical session border controller, all in only 9U of rack space. It supports a single, unique signaling interface per peering service provider, enterprise customer, or business/residential subscriber group. With the Net-Net PAC, a service provider can easily grow a single 1U Net-Net SD to a pack of eight, using the Net-Net SR for SIP-based load balancing to ensure optimum session performance and to protect against system failure.

2.1 Products and Versions Tested

The Net-Net PAC under test was a "half PAC" consisting of one Net-Net SR and four Net-Net SDs in both high-availability (HA) and non-HA configurations. The "half PAC" represented half the performance and capacity of the product described in the paragraph above. In the HA configuration, each active system had a standby system. All media and signaling state was check-pointed to the standby system that took over in the event of active system failure. The hardware for all ten systems was Part Number 721-1502-02 Rev. 1.03 and spanned Serial Numbers 010416003050 to 010416003059. Each system ran pre-beta Version 2.0 software without production compiler optimization. Using this software image, the systems were configured as Net-Net SRs or SDs.

3 Test Philosophy and Setup

Testing network devices that provide "mission critical" functions within an operating IP network delivering voice services is challenging. Networks and equipment do not always function "as expected," especially when under heavy load, with immature protocols, or when components fail. Network devices must be validated under conditions that as closely as possible mimic situations and scenarios that will be experienced in the real world.

CT Labs strongly believes that real world testing with transactionally accurate call flows with real media is the best way to accurately validate these types of devices. This places a heavy burden on both the design of the tests and the test equipment used. The real world conditions must be mimicked completely, with test equipment that can act and respond precisely like the real networks, network devices, and end users that will interface with the network devices under test. This requires that the test equipment maintain accurate and realistic call state. For example, if during a call a network device does not respond in time, the test system must "back off" like a real user and retry. Every effort was made to assure transactionally accurate and real world conditions during this test.

In accordance with this test philosophy, CT Labs utilized two VoIP test tools from Empirix for this test:

1. The **Hammer NXT™**, a scalable IP/TDM performance test platform, which generates complete and transactionally accurate calls with integrated signaling and media. It provided call volume verification, media path confirmation and voice quality scoring on all calls.

2. The **Hammer Call Analyzer™**, a call and protocol analysis tool that simplifies the debugging of VoIP networks, provided SIP call debug and analysis, real-time QoS metrics, Voice Quality scoring and a multistage graphical display for each call flow. CT Labs also used the call ladder tool to determine signaling latency for individual calls.

Figure 1 – Acme Packet Test Bed

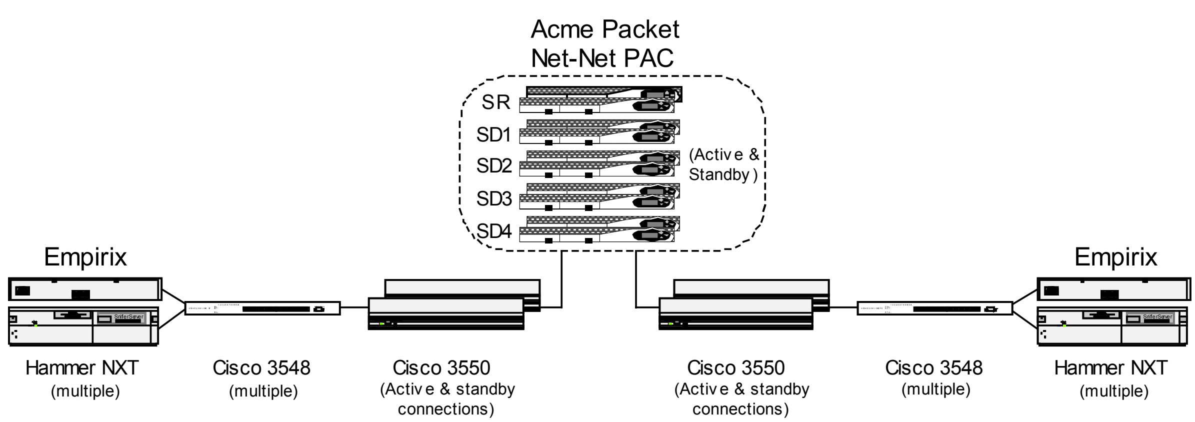

For all tests, 50,000 concurrent calls with full G.729 conversations were established across the Hammer NXT test setup. The total calls per second (CPS) rates were divided between the NXT call/answer test units. Gigabit fiber connections were used for all systems between the Cisco 3548 switches. The 3548s consolidated nine Fast Ethernet connections across each Hammer NXT test setup.

The Hammer NXT test scripts were designed to test full Network Address Port Translation (NAPT) at level 3 and 5. Once a call was established between calling and answering Hammer NXT units, G.729 encoded speech was sent for the duration of the call. For the maximum call rate discovery tests, the simple three-message set-up scripts (INV, OK, ACK) were used. For the real world call rate tests, the more typical seven-message set-up and tear-down scripts (INV, TRY, RING, OK, ACK, BYE, and OK) were used.

4 Testing Results

4.1 Results Highlights

The following is a summary of results for the tests conducted on the Net-Net "half" PAC:

1. A call rate of 450 call set-ups per second was achieved with media using Session Initiation Protocol (SIP) with complete NAPT at layers 3 and 5. This is equivalent to 1.62 million busy hour call attempts (BHCA).

2. A call rate of 200 calls per second was achieved with media for calls utilizing a more typical real world seven-message SIP call flow including continuous call set-up and tear-down (INV, TRY, RING, OK, ACK, BYE, and OK) with complete NAPT at layers 3 and 5.

3. A total of 50,000 concurrent active calls were supported in both tests. Every call integrated both signaling and media and was transactionally accurate.

4. All test runs were observed to execute without any loss of active calls and without any loss of ability to set-up or terminate calls when tested in the high availability configuration and subjected to individual Net-Net PAC Session Router or Session Director failures.

5. The signaling latency for all test runs was observed to be in the range of 1-8 milliseconds.

Note that every test run was executed after a clean restart of Acme Packet and Empirix equipment to ensure repeatable and verifiably clean results from each test run. Each NXT system was set to generate a target CPS; the actual measured CPS rates varied slightly from the target CPS and are listed in Table 1 in the next section. Because of the network configuration's total bandwidth limitations, for this test only G.729 media streams were used. It should be noted that the Acme Net-Net PAC is codec agnostic and supports all voice and video codecs, and that the Hammer test systems support all popular codecs.

4.2 Maximum Call Rate Discovery Test Results

To demonstrate the absolute maximum call-handling capacity of the Net-Net "half" PAC, a simple three-message set-up script was used (i.e. INV, OK, ACK). One set of results was produced using the high availability (HA) configuration and the other without HA as presented in Table 1 and Figure 2 below.

Table 1: Maximum Call Rate Discovery Performance Results

SR = Session Router SD = Session Director HA = High Availability

Performance Parameter	Non-HA Results	HA Results
Calls Per Second (CPS)		
- Maximum	476	306
- Average	452	291
Maximum SR CPU Utilization	27.5%	19.3%

CT Labs found that, as expected, the maximum CPU utilization events occurred during peak call ramp-up. Given that the test configuration was a "half PAC", the Net-Net SR CPU utilization data suggests that the Session Router has the CPU capacity to easily service twice as many Net-Net SDs to deliver a projected 900 CPS in non-HA and nearly 600 CPS in HA "full PAC" configurations.

During the HA test and shortly after all calls were established, the first Session Director was restarted via the reset button. CT Labs then verified that all calls already in progress were unaffected. Then the Session Router was also reset. Again, CT Labs verified that all calls already in progress were unaffected The reduced call rate for the HA test occurred due to the additional overhead of check-pointing both media and signaling state between active and standby Session Directors.

Figure 2: Maximum Call Rate Discovery CPS Results

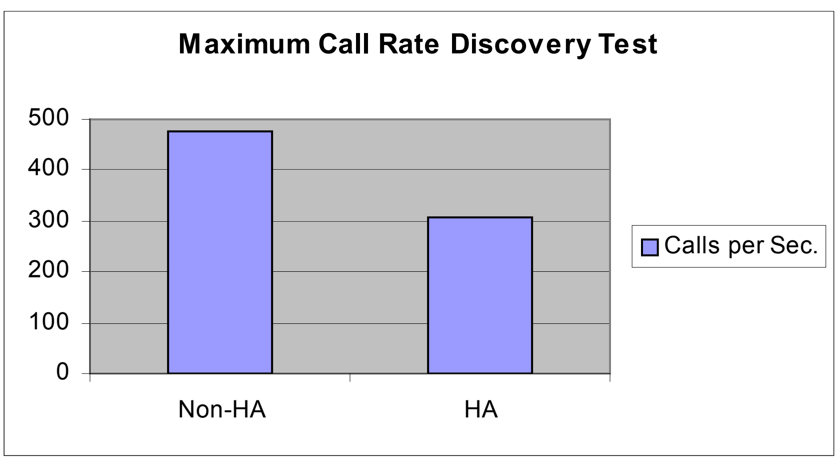

4.3 Real World Call Rate Test Results

To demonstrate the real world call-handling capacity of the product, seven-message set-up and tear-down scripts were used (i.e. INV, TRY, RING, OK, ACK, BYE, and OK). The Hammer NXT test platform generated unique, transactionally accurate and stateful signaling and media flows for every call as shown in Figure 3 below.

Figure 3: Real World SIP Call flow

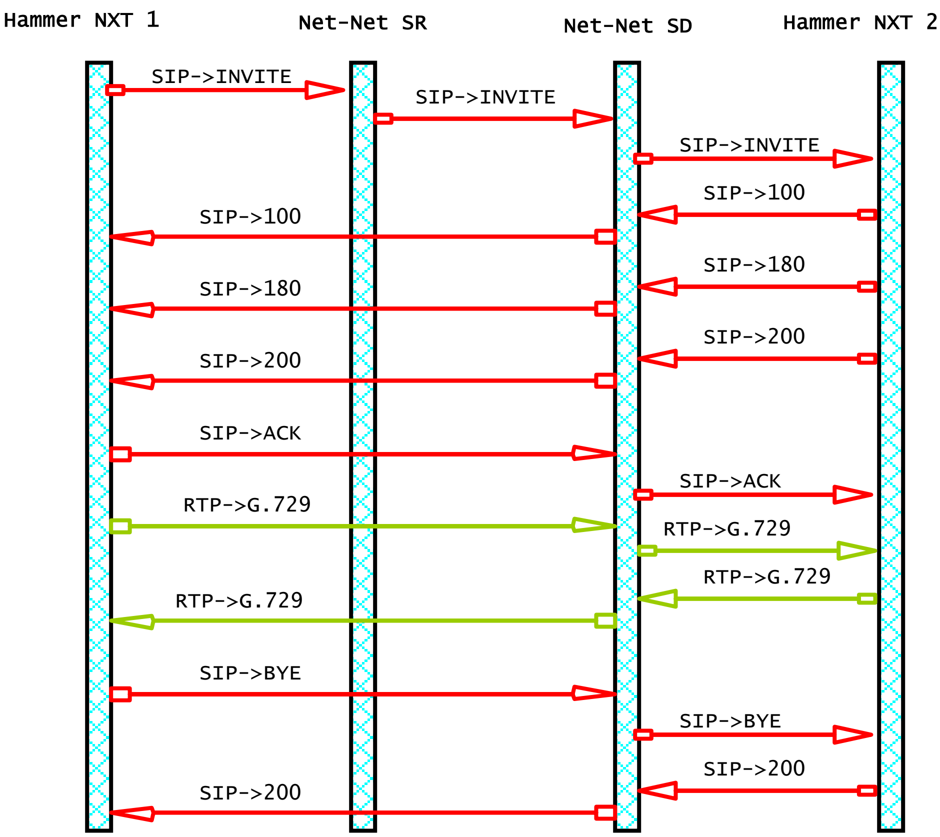

Two sets of results were produced using the high availability (HA) configuration, and the other using a configuration without HA. The results of these tests are presented in Table 2 and Figure 4 below.

Table 2: Real World Call Rate Performance Results

SR = Session Router SD = Session Director HA = High Availability

Performance Parameter	Non-HA Results	HA Results
Calls Per Second (CPS) – Average	202	120
Maximum SR CPU Utilization (%)	13.4%	8.2%

Again, the Net-Net SR utilization was found to be very low, suggesting that the Net-Net SR could handle many more Net-Net SDs than was used for this test. Given that the test configuration was a "half PAC", the Net-Net SR CPU utilization data suggests that the Session Router has the CPU capacity to easily service twice as many Net-Net SDs to deliver a projected 400 CPS in non-HA and 240 CPS in HA "full PAC" configurations for these real world call flows. The reduced call rate for the HA test (from 202 to 120 calls per second) occurred due to the additional overhead of check-pointing both media and signaling state between active and standby Session Directors.

Figure 4: Real World Call Rate Discovery CPS Results

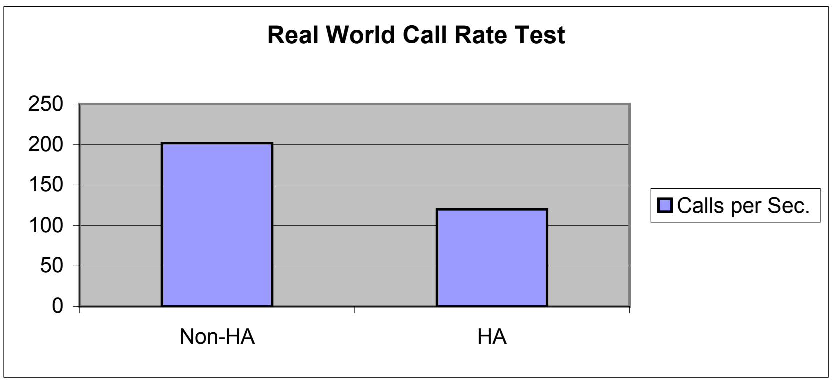

Two HA failover tests were run: one for Net-Net SR failure and another for Net-Net SD failure. This allowed separately tracking the effects of each type of failure. For the Net-Net SR failover the NXT signaling test platforms recorded 154 total call placement failures. Such failures are expected for this test because any INVITEs in the Net-Net SR at the time of failure can fail with a response timeout, and the 154 failures accounts for less than two seconds worth of calls. By the time that number of subscribers retried the calls, the system would be available for service again. For the Net-Net SD failover test, the network cable was disconnected from the active Net-Net SD while calls were ending and restarting. During this action, no calls were observed to be dropped and no failures of any kind were detected for the 30,000 calls that completed after the failover.

Signaling latency at a call rate of 100 CPS for the HA configuration was also measured. The signaling latency was observed to be in the range of 1-8 milliseconds depending upon SIP message as detailed in Table 3 below. Maximum SR CPU utilization was measured at 7.1%.

Table 3: Real World Call Flow – Signaling Latency Results

SR = Session Router SD = Session Director HA = High Availability

Performance Parameter	HA Results
INVITE SR-SD Signaling Latency[1] (mS)	
Minimum	7.36
Maximum	7.93
Average	7.66
INVITE SD Signaling Latency[2] (mS)	
Minimum	3.97
Maximum	5.37
Average	4.47
OK Signaling Latency (mS)	
Minimum	2.03
Maximum	4.79
Average	2.52
RINGING Signaling Latency (mS)	
Minimum	1.27
Maximum	1.48
Average	1.35

5 Conclusion

Through the tests performed during this project, CT Labs was able to confirm the high performance, availability and capacity of the Acme Packet Net-Net "half" PAC configuration for SIP-based calls, as shown in Figure 1 of this report.

CT LABS TESTING SERVICES

The Converged Communication Industry's Premier Testing Service

CT Labs testing solutions for converged communications equipment include automated real-world performance testing, user interface evaluations, and functional, regression, and load & stress tests. Managed conformance testing services are available for a wide range of communications products. The *CT Labs Tested* program provides independent validation of usability and product performance.

[1] *INVITE SR-SD Signaling Latency* measures the time that elapses from the ingress of the INVITE signal on the SR to the egress of the INVITE signal on the SD. In order to measure the worst-case impact, measurements were only taken during ramp up (peak load on the system).

[2] *INVITE SD, OK, and Ringing Signaling Latencies* measure the time that elapses from the ingress of the associated signal on the SD to the egress of that signaling packet from the same SD. In order to measure the worst-case impact, measurements were only taken during ramp up (peak load on the system).