

SIP and RTP Denial of Service Attack Tests Summary Report

CT Labs was commissioned by Acme Packet to verify that their Session Border Controller (SBC) product, the Net-Net Session Director, provided a robust level of security for service providers with respect to a wide range of potential attacks designed to degrade or terminate their services. To accomplish this, a test plan was created with the goal of verifying proper operation of the Session Director (SD) during a variety of denial-of-service (DoS), distributed-denial-of-service (DDoS), and other attacks on SIP-based services. An unprotected SIP Proxy Server and Firewall were also tested, to provide comparison data and highlight their vulnerability.

Statement of Test Purpose

The purpose of the tests was to evaluate the Acme Packet Net-Net Session Director session border controller and verify that it could effectively resist DoS and DDoS attacks without impacting established and new calls from valid / legitimate users. For the testing, SIP-based VoIP and RTP traffic was used to evaluate the impact of attacks on valid signaling and media data.

Products and Equipment Staged

Products Tested

- Acme Packet Net-Net SD platform, software version 2.0.
- SIP Express Router (SER), a high-performance, open-source SIP server available from www.iptel.org
- The highest end hardware and latest software of a leading third-party high performance firewall product (well-known vendor name withheld)

Test Equipment and Tools

- Gulp, version 3.0 is an Acme-created attack tool, to create malformed or well-formed packets and flood them as fast as possible to an attack target.
- Empirix Hammer FX-IP version 2.1, used to generate legitimate call traffic with RTP streams from 576 users.
- SIPp, a well known SIP client and server scriptable emulator, available from <http://sipp.sourceforge.net>
- Network protocol analyzer.



This summary report highlights test results for the SIP attack tests performed by CT Labs. Please contact Acme Packet for a copy of the full-length report.

Highlights

- *Acme Packet Session Director passes grueling CT Labs attack test with flying colors, zero call failures*
- *Acme passes 7.6 million SIP calls under attack of 40 billion fraudulent Invite messages, from over 1 billion random sources, in 60-hour reliability test*
- *Tests highlight fundamental architecture differences between SIP proxy, firewall, and session border controller products*

Executive Summary

CT Labs performed testing to measure the protection that the Net-Net SD session border controller would provide to a private network during a variety of SIP attack scenarios. One of the main criteria of the test was to ensure that legitimate SIP calls were not significantly affected while the Net-Net SD provided protection against SIP attack packets.

CT Labs found that the Net-Net SD provided excellent attack protection, as the attack tests did not measurably impact call performance of SIP VoIP traffic with the DoS/DDoS attack scenarios lodged against it. In all tests that measured SIP call performance, the average call setup latency did not increase by more than 16 ms, and for most, it did not increase by more than 2 ms. RTP media jitter for all test cases did not increase by any value measurable by the test equipment. This is an excellent result.

CT Labs further found the Acme Packet Net-Net Session Director product to perform flawlessly in our tests, not only exhibiting a negligible impact on performance but also succeeding in passing 100% of the intended SIP call traffic without dropping even a single call during the comprehensive series of attacks.

Testing Methodology

Three test topologies were used: one without a security device and the SIP Proxy as the DUT (Setup #1); one with a firewall as the DUT (Setup #2) protecting the SIP Proxy, and one with the Acme Net-Net SD as the DUT (Setup #3) protecting the SIP Proxy. Each topology included two sets of callers: a group of dynamically registered endpoints representing a consumer access service, and a group of statically provisioned devices representing peering partners. For this evaluation, tests were designed to simulate a variety of SIP and RTP attacks. In most tests, legitimate SIP calls using continuous RTP media streams were placed using the Empirix Hammer FX-IPs through the SD and a protected SER proxy server while various DoS or DDoS attacks were launched against the protected network.

Tests were designed to verify that:

- The Net-Net SD did not fail while under attack
- The SER proxy server was protected by the Net-Net SD
- Valid SIP VoIP calls by legitimate users were not impacted by the simultaneous attacks

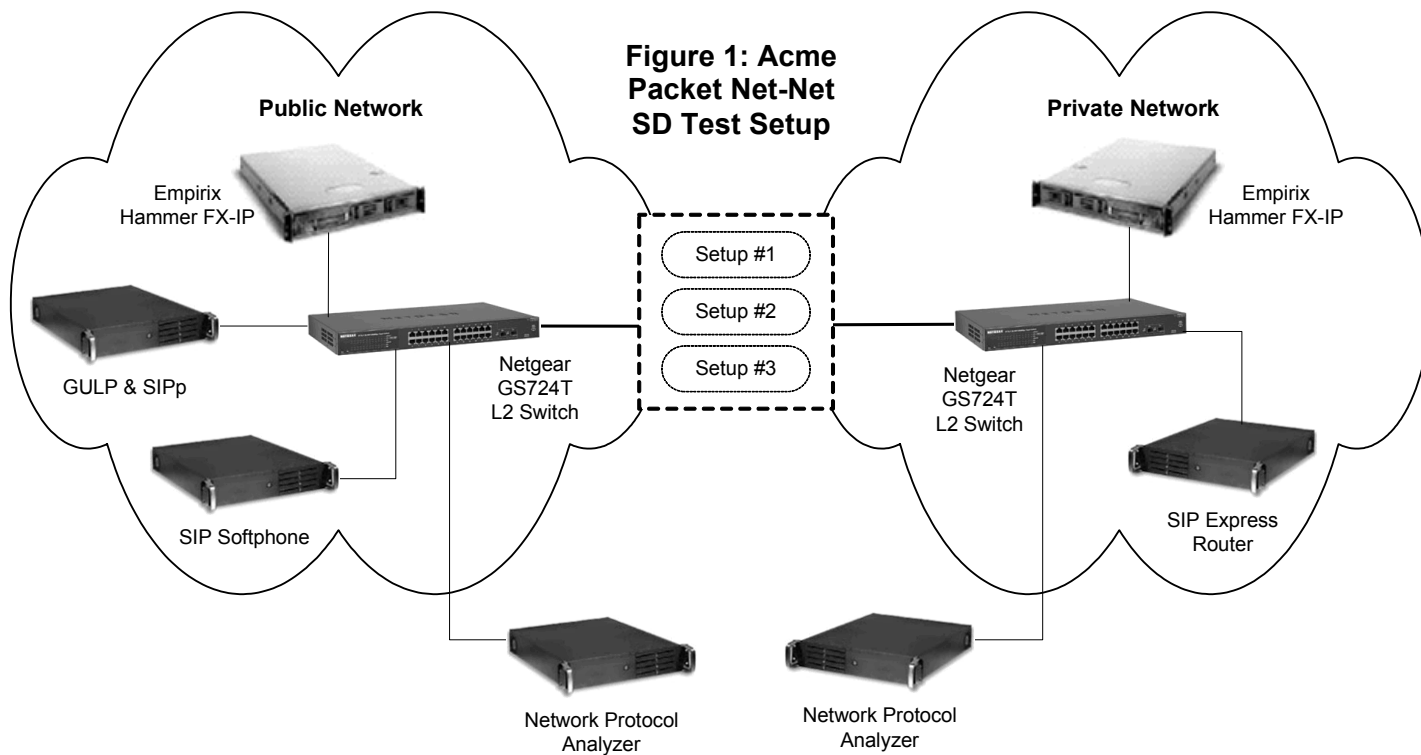
Baseline Performance Tests

Baseline tests were performed on the test setup shown in Figure 1 using the Empirix Hammer FX-IP VoIP call generators, to compare performance before and during attack tests. The baseline measured signaling and media latencies end-to-end, from one FX-IP, through the Ethernet switches, Net-Net SD, and Iptel SER to the other FX-IP.

Subsequent tests running DoS and DDoS attacks were then measured to determine the increase in signaling or media latency or jitter, if any.

Performance Test Results

As can be seen in the graph in Figure 2 on the next page, the DoS and DDoS attacks produced no significant increase in signaling delay or media jitter on the Net-Net SD. The average increase in signaling delay was within the measurement error of the test equipment. The worst-case increase for the most intensive DDoS test increased call setup latency by an average of 16 ms, an insignificant value. Generally, signaling latency can be as high as 400 ms without user-perceived impact.



Equipment was configured as shown in Figure 1 above. A pair of Empirix Hammer FX-IP call generators was used to generate the SIP VoIP call traffic. SIPP, Gulp, SIP Express Router, a SIP Softphone, and network protocol analyzers were installed on CT Labs servers for use in this test. SIPP was used to generate SIP signaling packets for RTP Fraud testing, and Gulp was used to generate the simulated DoS/DDoS attacks. To test both access and peering models with a single configuration, the FX-IP in the public network was configured to simulate both a peering host and group of access SIP users. **Setup #1** was the baseline “straight-through” test; **Setup #2** was implemented with the third-party high-performance firewall product, and **Setup #3** utilized the Acme SD product.

SIP & RTP Attack Tests

SIP Flood Tests

CT Labs used the Gulp tool to create floods of INVITE, Response, and REGISTER messages in both peering and access modes, while making legitimate calls. The SIP messages were flooded at 500Mbps from thousands of random source addresses / ports, constantly changing SIP headers to avoid detection.

SIP Spoof Flood Tests

CT Labs used the Gulp tool to simulate several types of spoofing attacks with INVITE, Response, and REGISTER messages in both peering and access modes. Different fields, headers, and addresses were spoofed. Tests were designed to verify either no impact, or limiting the impact to affecting users being spoofed.

SIP Malformed Packet Tests

CT Labs utilized the Gulp tool to send the complete set of over 4500 Protos¹ attack test cases, as well as several other types which are known to cause problems for SIP devices

SIP Torture Tests

The IETF has created a draft of 49 malformed or unusually formatted SIP messages which can cause problems for poorly-designed SIP parsers. CT Labs used the Gulp tool to send the IETF test cases.

RTP Attack Tests

CT Labs used SIPp and Gulp to perform a variety of RTP fraud and denial of service attacks.

¹ The Protos test suite was developed by the University of Oulu, in Finland (www.ee.oulu.fi).

SIP & RTP Attack Test Results

During these tests, CT Labs observed that the Net-Net SD never failed to protect the SIP Proxy, or legitimate user calls, media, and service. The SIP proxy, in a standalone configuration, did not crash but effectively reached processing exhaustion at a fraction of the default attack rate, and thus failed to provide service while under attack. The firewall likewise failed to provide protection, by allowing the flood to impact the SIP proxy, or in some cases by becoming unresponsive to any input.

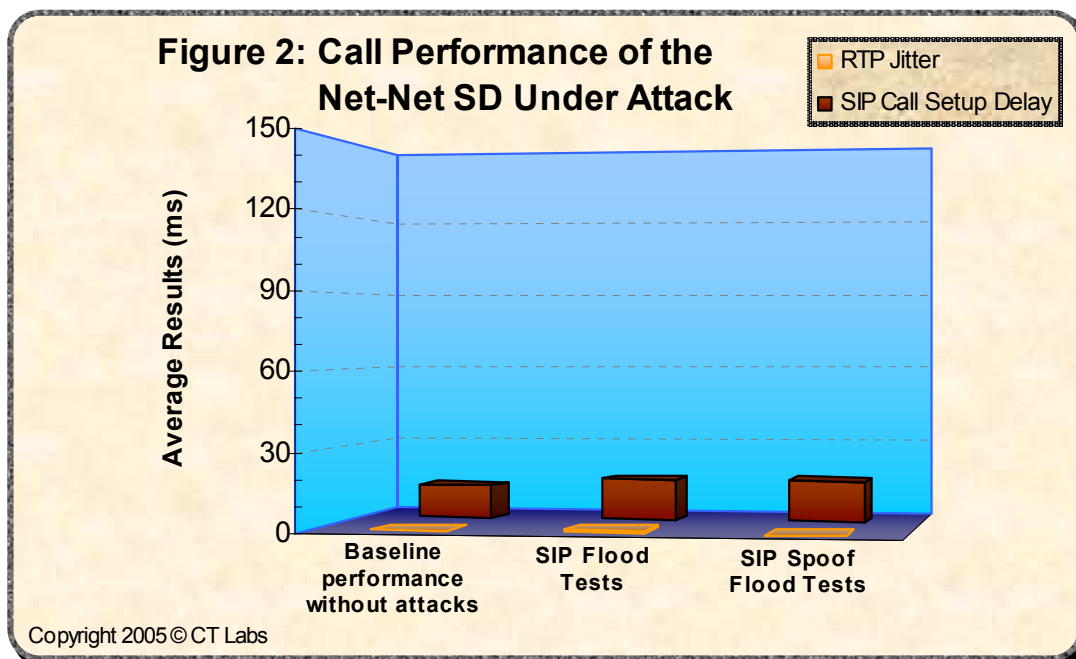
Proxy Server Comparison Test

This test was performed with the goal of determining the vulnerability of a high-performance but unprotected SIP Proxy server. The test setup used is shown in Figure 1, Setup #1 (Net-Net SD removed), allowing calls and attacks to be made directly to the SER proxy server.

Proxy Server Test Results

While the Iptel SER is an efficient proxy server, and was running on a high-end IBM server, ultimately it has no hardware protection for attack mitigation. Thus even simple DoS or DDoS flood tests easily consume sufficient CPU processing resources to cause legitimate user calls to fail.

CT Labs found after several short test runs that the Iptel SER was unable to handle DoS and DDoS tests, even when reducing the attack rate to a fraction of that used against the SD. For example, in a spoof flood test, even 1/10th the attack rate still overloaded the proxy. The SER did not crash, but its CPU load reached 100% preventing it from servicing new calls. RTP tests were not run, as the Iptel is a SIP signaling proxy: thus it does not handle media and does not provide any media protection.



Firewall Comparison Test

This test was performed with the goal of comparing the attack performance of the Net-Net SD to a high-end firewall product from a leading firewall vendor. The test setup used is shown in Figure 1, except that the Net-Net SD was replaced with the firewall product. While a SBC and a firewall are similar in that they are both designed to protect a trusted network from attacks launched from an un-trusted network, the capacities and services provided by each device are very different.

Firewall Test Results

CT Labs found after several short test runs that the firewall product was unable to handle the same capacity mix of SIP calls and attacks as was the Net-Net SD, nor provide the same fundamental type of protection. During this test the firewall product's performance was observed to either slow drastically or even halt, preventing further data from flowing through the device. Call quality through the firewall was found to quickly degrade with virtually any attack. In some cases the firewall would simply flood the attack through to the SIP Proxy, thereby causing an outage on the Proxy instead.

It should be noted that while a firewall is typically installed on a customer premise to protect a LAN, a SBC is typically installed in carrier facilities to isolate carrier networks. Thus, the results of this test highlight and confirm the fundamental differences of these devices and their ability to provide attack protection.

Reliability Test Run

Normally, obvious problems with the tested products were expected to surface within seconds after each test was initiated. However, each test case in this evaluation was run for at least 15 minutes to verify that there were no evident problems with these systems (i.e. table management errors, memory leaks). Additionally, some tests were performed for even longer periods of time to further evaluate the long-term stability and reliability of the Net-Net SD - including one continuous test run for 60 hours, and three runs for over 12 hours each.

Reliability Test Results

This 60 hour test of the Acme Packet Net-Net SD successfully completed 7.6 million SIP calls while simultaneously being attacked by a flood of over 40 billion fraudulent SIP Invite messages, from over 1 billion randomly generated source addresses from the entire IPv4 address range. The Invite messages were flooded at a rate of approximately 500Mbps, or 130,000 Invites/sec. No legitimate calls failed and no RTP media packets were lost during the 60 hour test. Three additional tests were performed of over 12 hours each, which similarly passed with no call failures or lost media. These included a flood of malformed, very small SIP Invite, Register and Options messages at approximately 300,000 messages/second.

Another notable long-duration test was conducted with flooding fraudulent Register messages from random sources, in an access scenario. This test is perhaps the most difficult to protect against, because a security device has to allow Register messages from any source in to the private network while not impacting legitimate calls and without impacting the SIP Proxy itself. CT Labs ran the test in the worst case scenario, whereby the source addresses/ports would constantly change so that the security device could not "learn" malicious devices over time. Even in this test, the Acme Packet SD allowed legitimate devices to make calls with no impact, and did not flood more Register messages to the SER proxy server.

About CT Labs

Background

CT Labs was founded in 1998 with the mission of providing outsource Q/A testing and marketing report services to the converged communications industry. The CT Labs team brings with it a wide range of talents and experience that gives us a unique ability to solve the most challenging test projects. Our open testing services philosophy enables us to provide our customers with test plans, test execution, testing reports, and even assistance in setting up specific testing environments in their own testing areas.

Facilities

Our test lab is well-equipped with tools and test platforms from our technology partners. In addition, CT Labs has the in-house expertise to develop specialized tools when off-the-shelf solutions are not available. CT Labs prides itself on keeping our lab current, enabling us to perform testing projects on cutting-edge next-generation network products and technologies.

www.ct-labs.com

v: 916-577-2100

f: 916-577-2101

info@ct-labs.com



Acme Packet, Net-Net Session Director and Session Aware Networking are trademarks of Acme Packet. Empirix, Hammer and Hammer FX are trademarks or registered trademarks of Empirix Inc. in the United States and other countries. All other trademarks contained herein are the property of their respective owners.