

## RE: Realsecure

*Source:* <http://www.derkeiler.com/Mailing-Lists/securityfocus/focus-ids/2001-10/0097.html>

---

**From:** Kevin Brown ([kbrownfox@home.com](mailto:kbrownfox@home.com))

**Date:** 10/16/01

From: "Kevin Brown" <[kbrownfox@home.com](mailto:kbrownfox@home.com)>  
To: "Jeff Nathan" <[jeff@wwti.com](mailto:jeff@wwti.com)>, "Greg Shipley" <[gshipley@neohapsis.com](mailto:gshipley@neohapsis.com)>  
Subject: RE: Realsecure  
Date: Mon, 15 Oct 2001 23:41:49 -0400  
Message-ID: <NEBBIFJCILKEKMJKGMFIKEAJCPAA.kbrownfox@home.com>

We have in fact been working with several vendors to better develop a better test methodology for IDS performance. Jeff is right, Chariot doesn't really scale well beyond 100 Mbps. We are developing a test methodology that would break performance into 3 separate scenarios, packets per second, megabits per second, and TCP sessions per second. The goal would be to isolate each of these metrics separately to determine which are the limiting factors for each individual product.

The first test would be done using lots of small, empty, layer 3 packets. By only building a layer 3 packet, we can eliminate variables such as port numbers while also limiting the impact of Mbps. The second test would be to generate large, complete 7 layer packets with as few connections as possible. And our third test would just test TCP connections to a server, both connections per second and total number of connections. So far, the response from the vendor community has been positive.

I believe that these types of tests coupled with the type of test Shipley conducted earlier this year would give readers a more complete picture of what to expect out of NIDS. I would never encourage anyone to trust only lab results any more than I would tell them to only trust "real world" tests. Both are needed to gain a complete picture of a product. I also believe that with these numbers, users could analyze the traffic on their own network and have a better idea of what to expect when they put these things in their own network. That is the end goal after all, isn't it?

Kevin D. Brown  
Miercom

-----Original Message-----

From: [jeff@dolavimus.vaughan.nu](mailto:jeff@dolavimus.vaughan.nu) [mailto:[jeff@dolavimus.vaughan.nu](mailto:jeff@dolavimus.vaughan.nu)] On Behalf Of Jeff Nathan  
Sent: Monday, October 15, 2001 2:26 AM  
To: Greg Shipley  
Cc: [focus-ids@securityfocus.com](mailto:focus-ids@securityfocus.com); Bob Walder  
Subject: Re: Realsecure

RE: Realsecure

## SecurityFocus IDS: RE: Realsecure

It's been said before and I'll say it again. We need a common framework for testing network ID systems. The NetworkWorldFusion report was just as Kevin Brown mentioned, thin. The attacks used were IMHO, neither extensive nor particularly representative of what we're seeing out there in the large enterprises.

As Greg and likely Dragos will attest, generating high speed traffic with NetIQ (formerly Ganymede) Chariot is just not possible. Unfortunately, the way many people test ID systems is in a lab with either canned or generated background traffic. This is not representative of real traffic. Neither are the number (and frequency of) of TCP sessions (and end-to-end UDP sessions if the ID system attempts to track UDP state ala NFR) representative of what exists on a real network – thus even beginning to test the state functionality and stream reassembly functionality of the ID system.

The only way to test an ID system is with real traffic. Just as others have done in the past, I will argue this issue vehemently. I've used Ixia and Smartbits devices as well as Chariot and I've looked at the traffic they generate. The traffic these systems generate serves only to load small portions of the overall ID system and while this conversation has come up in the past and been argued from many sides, only until the NWC test did we actually see ID systems really tested.

We can use hardware traffic generators to create different Ethernet frame sizes until the end of time or we can step up to the plate – when publishing – and perform tests that use real enterprise networks to test the systems. That way, the traffic validates itself.

There are a number of other technical questions these reviews haven't answer and will never answer: where did the system actually break down? When we talk about Ethernet frame sizes for example, how does one know if the performance of the ID system is completely hindered due to an Ethernet card with either too few or too small buffers? If this is the case, how do we know a particular system wouldn't have done better with another Ethernet card?

–Jeff

- 
- ***Previous message:*** [Reeves, Michael \(GEAE, Compaq\): "RE: Realsecure"](#)
  - ***In reply to:*** [Jeff Nathan: "Re: Realsecure"](#)
  - ***Next in thread:*** [Carric Dooley: "RE: Realsecure"](#)
  - ***Next in thread:*** [malj31: "Re: Realsecure"](#)
  - ***Reply:*** [Carric Dooley: "RE: Realsecure"](#)
  - ***Messages sorted by:*** [\[ date \] \[ thread \] \[ subject \] \[ author \] \[ attachment \]](#)