

Internet Leukocytes Need Your Help

An ecosystem is defined as all the members of a biological community and the physical environment, functioning as a unit (Answers.com, 2008). Nutrients move around the ecosystem in loops ensuring that all members are properly nourished. All ecosystems are open systems in the sense that energy and matter are transferred in and out. The web of life interconnects all aspects of the ecosystem needing to maintain a balance in order to fully function and survive.

The Internet could be seen as one such ecosystem only it is an ecosystem that is currently ill with certain habitats that exist exclusively to act as parasites, not to destroy the ecosystem but to continuously feed off the weak in some Darwinian manner. The concept of an immune system




Figure 1 - Source caida.org

for the Internet is then a requirement for sustainability:

“The biological immune system provides a compelling example of a massively-parallel adaptive information processing system, one which we can study for the purpose of designing better artificial systems. It exhibits many

properties that we would like to

incorporate into artificial systems: It is diverse, distributed, error tolerant, dynamic, self-monitoring (or self-aware) and adaptable. These properties give the immune system certain key characteristics that most artificial systems today lack: robustness, adaptivity and autonomy.” (Forrest and Hofmeyer, 2000)

If the Internet was in fact a self-sustaining ecosystem with a built in biological immune system, then self-healing could occur and safety could be assured. But there is one particular factor that skews the innate immune system responses that are triggered by malicious activities or toxins to the ecosystem – humans.

DNS AND CACHE POISONING



Figure 2 Source computerworld.com.au

Take for example the issues associated with DNS Open Recursion. There continues to be a significant Internet threat due to the domain name system (DNS) Open Recursion capability that is enabled by default in many DNS implementations. Implementations that do not follow proper configuration of the DNS. If the DNS was properly *configured* by humans and configured by humans to be less trusting, the issues would largely go away yet the introduction of human interaction into the ecosystem guarantees inherent weaknesses that befuddle the innate immune system of the Internet.

Configuration management focuses on establishing and maintaining consistency of components within the ecosystem ensuring proper performance of functional and physical attributes, and most importantly, ensuring malicious activity cannot injure or weaken the ecosystem. Of course, the configuration comes from the vendor (humans) (in most cases) in an open and vulnerable state. Why deliver a product that is exploitable by known threats out of the box?

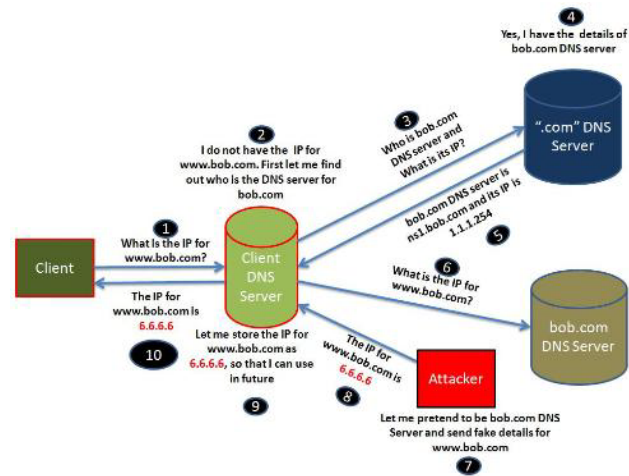
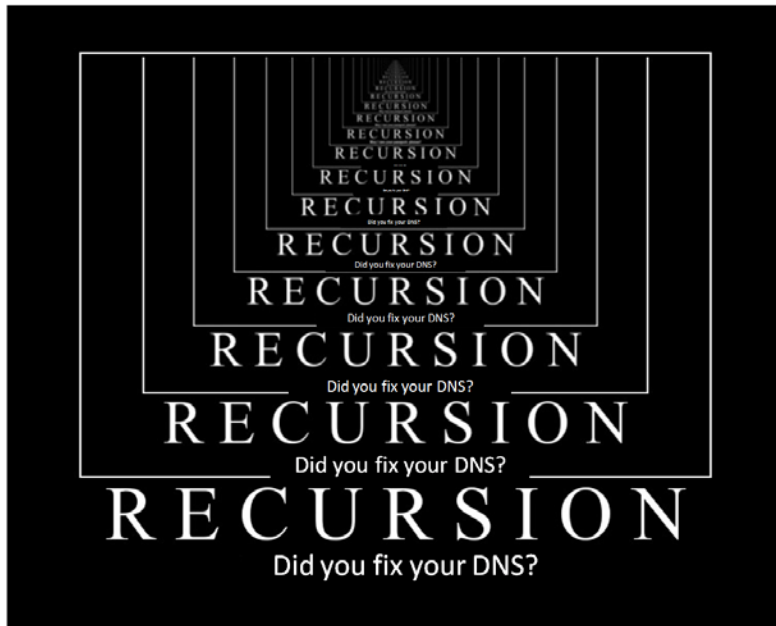


Figure 3 Source vil.nai.com

The DNS is the component of network infrastructure that maps human-readable labels (names) to Internet addresses and other resources. The two potential threats are:

1. Cache poisoning which occurs when malicious or misleading data finds its way into a DNS caching server. The bad data is then made available to programs running on workstations (e.g., web browsers and email applications) that request the cached data. These applications are then redirected to send data to compromised hosts or black holes.
2. Distributed denial of service (DDoS) attacks could take out most any Internet site including generic top-level domain (gTLD) name servers (e.g., .com, .net, .org). An



attacker could launch source address spoofed stream of queries, which could result in an amplified reflection attack directed back at the spoofed victim address.

Root DNS servers do not allow this recursion feature; top-level domain (TLD) servers generally do not allow recursion. Configurations for servers below the TLDs are determined by each administrator

or user, but in reports is widely available on the Internet. This succinct description of the problem is further explained as documented by the CERT/CC in [VU109475](https://www.cert.org/advisories/CA-2003-07).

Extensive private sector research shows many name servers are configured to accept and perform recursive queries for anyone. This means that anyone can launch a query to a name server for any resource record. This can potentially be solved through proper configuration.

All organizations (yes – humans again) should determine if their domain servers allow external recursion and evaluate the requirement to allow recursive queries on their internal name servers. A single solution (e.g., script/fix/patch) is not possible due to the immense number of internal/external DNS configurations. While on the surface, this appears a daunting task, it is a rather simple procedure that should take system administrators an hour or two per domain server to accomplish.

The following courses of action are recommended to protect name servers:

- The system administrator restricts recursive lookups only to the authoritative domains served by the DNS server and does not allow world recursive lookups for the server.
 - This method helps ensure the DNS server's cache does not become poisoned for another domain and that this server does not participate in reflection attacks.

ADDRESS SPOOFING

- Log checking scripts that monitor the overall activity of the name server and look for DNS lookups from outside the domain alert the administrator to an outsider possibly looking for a DNS server vulnerable to cache poisoning.
- For those organizations operating internal and external name servers, authoritative and recursive service should be separated. Basically, two separate DNS server types are maintained, with two separate sets of data.
 - Internal DNS servers should only accept requests from internal hosts.
 - The external DNS server only accepts Queries for authoritative names the DNS server maintains.



This is but one example of what happens when the human factor is introduced into an ecosystem. If the FBI had a top ten list of threats to the Internet, the DNS Open Recursion problem would be near the top. Now that a solution to the problem is available, why not update the immune system with the appropriate fix and remove the ability for exploitation? Okay, okay you are still not convinced you

have the time to check your DNS servers to determine if they are vulnerable to these attacks.

Well, one such organization is ready and able to help. Simply email Team Cymru at info@cymru.com if you would like them to send you a list of your publicly accessible DNS servers that are capable of participating in one of these attacks. This information, coupled with the reference links below will hopefully help alert you to any potential open resolvers in your network, and provide some tips on how to fix them.

Read more details of the project at:

<http://www.team-cymru.org/Services/Resolvers/>

See the latest "Who and Why Show" which explains the problem

at <http://www.youtube.com/teamcymru>

Read the new white paper on this topic at:

<http://www.team-cymru.org/ReadingRoom/Whitepapers/>