Inter-domain Routing Stability Dynamics During Infrastructure Stress Events: The Internet Worm Menace

Francesco Palmieri
Universita degli Studi di Napoli Federico II, Centro Servizi Didattico Scientifico
Via Cinthia 45, 80126, Napoli, Italy (Email: fpalmieri@unina.it)
(Received Dec. 31, 2005; revised and accepted Apr. 25, 2006)

Abstract

The Internet is crucial to business, government, education and many other facets of society and its continuing scalability places serious challenges on the routing system's capability to produce a stable view of the overall network reachability. Several global-scale Internet failures driven by the uncontrollable spreading of self-propagating code exploiting homogeneous security vulnerabilities have led the popular press to predict the imminent death of the Internet. The last few years have seen a dramatic increase in the frequency and virulence of such "internet worm" outbreaks infecting hundreds of thousands of Internet hosts in a very short period and sometimes disrupting the connectivity of some large sections of the Internet with significant damages in network stability. Although the predicted Internet col-lapse has yet to materialize, further analysis of the behavior and characteristics of wide-area network dynamics during these events is critical for the evolution of the Internet, and there is considerable effort focused on developing technical means for detecting and containing worm infections before they can affect global network stability. The main focus of this work is the impact of worm spreading events on BGP inter-domain routing dynamics on time scales that are long enough to have the potential to increase route convergence times and impact the network behavior. Our analysis shows the direct correlation between observed instabilities and some worm attacks. We analyzed in detail the Internet infrastructure failure and pathological events that can be triggered by the abnormal traffic during the worm attack and how they can lead to global routing instability.

Keywords: BGP, inter-domain routing stability, worms

1 Introduction

In a few years, the Internet has evolved from a relatively obscure, experimental research and academic network to a commodity, mission-critical component of the public telecommunication infrastructure. Internet backbone failures that previously only impacted a handful of academic re-searchers and computer scientists, may now as easily generate millions of dollars of losses in e-commerce revenue and interrupt the daily routine of hundreds of thousands end-users. Unfortunately, from the security and survivability point of view, the Internet today can still be compared to the Wild West of the late 19th century United States. While this analogy may go a bit far, it is very important, however, to have an effective way of studying various aspects of the Internet behavior to improve its infrastructural security and stability. Today, the Internet connectivity can be disrupted by link or equipment failure, unintentional errors or, more frequently, malicious infrastructure-wide attacks, which have happened on several occasions. Typically, these infra-structure stress events can damage communication on many critical inter-domain links and have a direct impact on rout-ing stability and hence on the overall network functionality. In the last few years we observed that malicious computer-related stress events, such as the massive worldwide spreading of very fast and aggressive worms, may almost completely disrupt the connectivity of large sections of the Internet, affecting inter-domain routing behavior and severely damaging the network stability like and more than severe physical failures due to catastrophic events. This is essentially due to the fact that the easy access and wide us-age of the Internet make it a primary target for the propagation of worms. Since the appearance of the first well-known Internet worm, known as the Morris worm [11], in 1988, attackers have continuously developed more and more dangerous and fast-spreading worms. Today, our computing infrastructure is more vulnerable [7] than ever before since the relatively homogeneous software base in almost all the networked nodes coupled with the current high-bandwidth connectivity provides an ideal climate for self-propagating attacks. One might expect that as effective defenses are erected to protect against each new worm threat, these threats would be eliminated. The reality is quite different, a large number of threats never really die off. Data collected on Internet worms using a network of distributed blackhole sen-

sors monitoring blocks of unused address space over a period of a month and a half showed 75K Blaster worm infected hosts, 180K Slammer worm infected hosts, and 55K Code Red II worm hosts [1]. These numbers are quite surprising considering the Slammer worm and the Blaster worm were released nearly two year ago and the Code Red II worm over three years ago. The persistent presence of these long-lived worms illustrates how these malicious agents may become a consistent menace for today's Internet survivability. Of course the direct consequences are the enormous operational expenses to track down, contain, and repair each infected machine. In response to this threat, there is considerable effort focused on developing technical means for detecting and containing worm infections before they can cause any damage. The focus of this paper is on the emergence of noticeable inter-domain routing instability events in the Internet that can be observed in the collective behavior of the BGP routing traffic dynamics during the most infamous worm outbreaks happened in the last years. The BGP message streams and statistics collected by the Oregon Route-Views and RIPE NCC monitoring points during each month in which a worldwide worm infection has been happened, have been analyzed and correlated at multiple resolutions. In association with all the most aggressive worm infections we detected hours-long periods of exponential growth and decay in the route change and pre-fix announcement/withdrawal rates across all the observed peering sessions, indicating significant widespread degradation in the end-to-end functionality of the global Internet. We have documented so far a compelling connection between these global inter-domain routing instabilities and the propagation phase of the fastest-spreading Microsoft worms such as Code Red, Nimda and Slammer.

2 Internet Worms

Although the first well-known replicating program (the so-called Morris worm [10]) that self-propagated across the Internet by exploiting security vulnerabilities in host software was observed many years ago, in 1988, only later in 2001, Code Red, Code Red II, and Nimda showed us how our networks were still vulnerable. They infected a vary large number of hosts in the Internet in a few days, causing millions-of-dollar loss to our society and originating a new wave of global-scale propagating worms. On January 2003, Slammer was released and quickly spread throughout the Internet [6]. Because of its super fast scan rate, Slammer infected more than 90% of the vulnerable computers on the Internet within 10 minutes [6]. In addition, the large amount of scan packets sent out by Slammer caused a global-scale denial of service attack to the Internet; many networks across Europe, Asia and America were effectively shut down for several hours. Only a half-year later, the Blaster worm appeared and infected more than 200,000 computers within a couple of hours on August 2003. Finally, on the first days of May 2004 a new very noxious worm, known as Sasser has been found spreading in the wild and infecting millions of Windows-based host, affecting large section of the Internet. Code Red, Code Red II, Nimda, Slammer and Sasser have created a new generation of global-scale fast spreading worms. All of them are scan-based worms that find and infect target machines by testing any IP address in the entire Internet address space. Some of them deployed several distinct spreading mechanisms to make the containment more and more difficult. As a consequence, to cope with the inevitable emergence of significantly more aggressive worms, really effective early detection strategies, based on the clever observation and deep understanding of network instability phenomena are required to trigger as soon as possible the proper infection containment countermeasures.

3 The Measurement Methodology

Many successful academic and commercial projects use direct link/traffic-based measurements (such as ping, traceroute, and web page access statistics data) to study the behavior and dynamics of the Internet under particular traffic load conditions. Such efforts are inherently limited by the locations of probe points required to "cover" the Internet meaningfully. This approach presents several visibility drawbacks, since simply placing agents throughout the Internet's core, as done by several commercial services, only builds up a picture of core-to-core traffic latencies and losses that has no power to predict the true "Internet weather" that end users really experience at the network edge. On the other side, studying the overall Internet BGP dynamic inter-domain routing data behavior from some well connected point in the Internet that has a complete knowledge of the Internet routing table gathered from multiple tier-1 peering associations, provides one of the few interesting alternatives to traffic-based analysis of the Inter-net's dynamics. This approach, while based on a more indirect observation of the network behavior doesn't suffer of the localization-dependent view problems of the traffic-based measurement model and will give, at our advise, a much better and globally scoped picture of the complex Internet dynamics that can be used to detect in an easy and reliable way all the infrastructure instability events triggered by widespread worm propagation.

3.1 Routing-Based Analysis of Internet Dynamics

The Internet is a large collection of Autonomous Systems (AS) or routing domains, which define separately administered networks that come under the control of a common authority, often corresponding to a network service provider. A little number of very large network service providers dominate the Internet. These international providers, often referred to as tier one providers, account for majority of resources and bandwidth that comprise

the public Internet. Several thousand of smaller regional networks, or tier two providers peer with the tier one providers in one or more large private or public Internet eXchange Points (IXPs). These exchange points are often considered the core of the Internet where national and international providers peer, or exchange both reachability information and traffic. There are various glues that hold together this massive system, but the most important of these is the inter-domain routing infrastructure. The Internet routing protocols maintain connectivity between and within AS's, and are designed to automatically reconfigure and recompute routing tables when they detect a link failure. This computation starts locally around the failure point, and then the information propagates through the Internet. The whole Internet routing infrastructure relies on the BGP protocol [8] to provide essential routing information between thousand of Autonomous Systems interconnected according to a multitude of business agreements. Every autonomous system connected to the Internet learns information about other networks by exchanging route information with its neighbors. An AS might have a single neighbor, such as an upstream ISP, or hundreds of neighbors in the case of tier-1 transit (backbone) networks, such as UUNET, Sprint, and Global Crossing, and the BGP protocol lets all the AS neighbors share information about network reachability. BGP conversations between border routers consist of two kinds of routing information messages: announcements and withdrawals. An announcement is a network reachability message, which specifies networks (identified as blocks of IP addresses) that are reachable via a particular route. Very often, a router changes its preferred route to a particular network and issues a new announcement. A withdrawal is a notification that a network is no longer reachable by any route via that neighbor. Border routers build a BGP routing table, or route information base (RIB), from accumulated BGP announcements and withdrawals from all neighbors. Depending on the number of neighbors it has, a border router could have several routing options for each destination network. From its BGP table, the border router chooses its preferred route based either on a selection algorithm or the routing policy specified by the network administrator. The number of networks in the table increases as the Internet grows and, today, a border router with a global routing table has about 500000 to 800000 networks in its RIB. Because the Internet depends exclusively on BGP for inter-domain routing, BGP-related failures and instabilities reflect immediately the worldwide network dynamics, so that we can derive performance, health and stability metrics by listening to "conversations" between routers, and specifically analyzing the BGP messages. By the very nature of globally distributed BGP routing processes, a listener located at any well-connected point, typically an Internet Exchange Point, has the opportunity to obtain a very accurate picture of the evolution of the overall reachability information in terms of best routes to every prefix in the Internet, delayed only by seconds to minutes.

3.2 Internet Routing Stability Metrics

The Internet continues to grow both in terms of its size and in terms of the services running on it. In addition, its continuing scalability places challenges on the routing system's capability to produce a stable view of the overall reachability of the Internet. Inter-domain routing instability, can be informally defined as the continuous, and sometimes rapid, change of network reachability and topology information such that it cannot be predictable and controlled by network Administrators. There are two predominant strategies for using BGP routing statistics to measure Internet instability. The first is based on the study of global reachability; that is, measuring the number of prefixes that appear in the Internet routing table at a given time. The second is based on the analysis of the inter-AS routing information rates of change; that is, measuring the number of prefix announcements and withdrawals in BGP UPDATE messages sent out and received on a well-connected BGP peering point per unit of time. According to both strategies, that can be easily combined to obtain a more effective picture of the Internet stability under stress conditions, the most interesting metrics that should be taken into account are:

- *BGP table size:* The BGP table contains "best routes" to all the destinations it has learned. Given a continuous BGP stream, starting at some time t in the past, sent from a neighboring router R, we can construct that portion of R's BGP table that has changed since time t. If a BGP session reset has occurred since t, then we can reconstruct the entire BGP table (since a router must send its entire table when a session is reestablished after a reset). This table can be thought of as a materialized view of the BGP update stream. One simple query over this view is the number of routes in the table, which can give us an approximate measure of the overall Internet reachability in term of announced prefixes. Prefix reachability doesn't necessarily mean network reachability, especially for the reason that a specific prefix can be contained within an aggregated route. However, an aggregated route is usually statically configured in the origin AS, and will never be withdrawn from the origin AS even though the more specific prefixes are no longer reachable in the origin AS. In this sense, the reachability of a specific prefix more realistically represents the reachability of the network than an aggregated one. Typically, during a worm attack affecting the network infrastructure stability we may note a sudden sharp variation in the number of routes in the BGP tables of the observed routers. When this phenomenon is observer across a large number of sources, it can be interpreted as a large-scale disruption in the global routing system and could be seen as an early warning sign of trouble on the Internet.

- *Update rate*: If we see large increases in the number of BGP update messages in a specific time interval, it's an unambiguous sign that the diversity of network prefixes is rising. Furthermore, the duration of these BGP message traffic surges, and the rate of their growth are what distinguish global instabilities from the pervasive background noise, daily rhythms and localized failures. Thus, from the total number and rate of BGP announcements and withdrawals for prefixes originated from an AS, it's possible to infer the stability status for the AS. The typical impact of a worm infection on BGP routing is easily observable as a sudden sharp increase in the number of BGP updates, primarily due to saturation of links at the edge of the network resulting in dropped BGP messages, which in turn causes BGP sessions to reset and BGP speaking routers to withdraw routes. In general, ISPs are very interested in monitoring the stability of routing updates sent from their customers since this can be the first symptom of a severe network failure.

- *Damping rate*: A more sophisticated measure of route instability can be based on the BGP's route flap damping mechanism [14], that was introduced specifically to prevent edge instability from flooding update messages globally. In detail, the "route flapping" phenomenon consisting in the rapid withdrawal and reannouncement of a route or a set of routes, typically due to a pathological condition on a BGP speaker, if not properly handled causes negative repercussions throughout the whole Internet. This nonlocality of effects is highly undesirable, and it is absolutely recommended to keep such effects naturally limited to a small area of the network around the problem. Accordingly the *BGP damping* feature explicitly inhibits a BGP router from exchanging "too many" messages about a given prefix with a given peer. Each route/prefix is associated with a penalty that is increased each time the route changes and once the penalty exceeds a threshold value, all the updates regarding the route are suppressed. This penalty decays exponentially using a configured halflife and affected routes are again accepted after the penalty falls below a re-use limit. The BGP route flap damping standard [14] recommends deployment only at "core routers", since when a flapping source is attached to a non-core node, the damping strategy is typically not effective in reducing the number of updates, and the deployment of damping will however imply a cost of delayed convergence. Of course, we expect a sharp increase in the number of "damped" routes during a worm infection that originates a noticeable infrastructure stress event causing route flaps, indicating a significant networkwide instability. Such change can be used as an alarm helping network operators to quickly recognize the problem.

The duration of these BGP phenomena, and the nature of their growth (linear or exponential, for example) are what distinguish truly global Internet instability from simple background noise (which is pervasive). Very short, high spikes in announcement or damping rates are very common since whenever a peer BGP session undergoes an hard reset, for example, a full table dump will follow. Furthermore, repeated session resets, due to fluctuations or reconfigurations in peering links of large transit ASes, may originate high and shortly-lived spikes in the overall BGP damping or update rate. More surprisingly, semi-permanent (typically some hours long) failures in the core internet infrastructure (fiber cuts, generator failures, building collapses, etc.) tend to generate only short-term increases in the BGP prefix announcement or BGP damping rates, which revert to the mean in a matter of seconds or minutes, as the highly redundant Internet core topology routes around the damage. Specific networks may remain unreachable until the damage is repaired, but because content networks are so vastly outnumbered by access networks, the "average" network prefix presumably adds very little in the way of marginal utility to the "average" Internet user. Of far greater concern are the appearances of sustained exponential rises in BGP message or event rates that last for hours. That is what the worm-triggered traffic causes.

4 The Stability Analysis

Our study uses BGP update data collected by The Oregon Route Views [12] and RIPE RIS [9] projects, both providing a service for Internet operators to obtain real-time information about the global routing system from the perspectives of several different locations around the Internet by peering with a wide variety of Autonomous Systems. The Route Views BGP monitoring project collects, since January 1999, update streams from more than 65 BGP speaking neighbors. On the other side the RIPE RIS project maintains several such monitoring points across Europe such as RIPE NCC and AMS-IX, (both in Amsterdam), LINX (London), SFINX (Paris), CIXP (Geneva), and VIX (Vienna). They peer with many of the so-called global tier-1 providers, plus very many smaller regional European networks. Access to multiple BGP monitoring points provides opportunities to filter the localized infrastructure failures that are close to individual collection points, clearing the way to unambiguously identify and study routing instability features that affect large portions of the Internet simultaneously. Typically, both projects collects, without providing any transit service, several million updates per day. This is approximately the same order of magnitude of updates that might be collected inside a large IXP with BGP streams arriving from many huge regions. Each AS router treats the monitoring point as a BGP peer router and sends routing updates to the monitoring point, where they are logged and made available to researchers. Therefore, by

establishing BGP peering sessions with a large number of BGP routers from well-connected organizations, analysis of traffic gathered at a single BGP monitoring point can provide a great deal of information about the way those organizations view the Internet, and about the dynamics of how paths change over a wide range of time scales. Of course, both the raw BGP message archives contain various errors and anomalies that require certain amount of care in their analysis, in particular: occasional timestamp clock shifts, missing data, corrupted message headers, truncated BGP messages, and several false extreme instability conditions due to the opening and closing of collecting routers' BGP sessions. Consequently, some preprocessing of the raw data has been performed to remove some of the above artifacts, such as instantaneous and short-lived announcement or withdraw spikes, that could else easily distort the results of our observations. Furthermore, since we were interested in long-lived instability events, to suppress the above misleading phenomena we smoothed the observation data by averaging them on a three hours sliding window and transformed them in absolute value differences on their monthly mean value. We operationally defined a noticeable global routing instability in terms of its rate, duration and diversity as follows: Exponential or similarly fast growth of the rate of prefix updates or damping rates, lasting hours to days (3 hours minimum), with a very large number of prefixes/networks churning (we used a cutting threshold set to the 70% of the maximum observed value).

4.1 The Observation Details

In our global routing stability analysis we considered the most noticeable worm infection events affecting the Internet in the last four years, as reported in CERT Coordination Center [2] Advisories or Incident Notes and in US-CERT Vulnerability Notes [13]. To know the human observed network behavior during these phenomena, we relied on case studies, news and reports from many network operators and organizations available on the Internet, as well as on the equipment vendor advisories. The detailed list of worm spreading events used in our analysis is reported below:

- *July 19, 2001 - Code Red (CA-2001-19)*
  The "Code Red" self-replicating malicious code exploits a known vulnerability in Microsoft IIS servers. The worm attempts to connect to TCP port 80 on a randomly chosen host and sends a crafted HTTP GET request to the victim, attempting to exploit a known buffer overflow in the Indexing Service. The same exploit is sent to each of the randomly chosen hosts due to the self-propagating nature of the worm. If the exploit is successful, the worm begins executing on the victim host and repeats its replication/infection cycle.

- *Aug 6, 2001 - Code Red II (IN-2001-09)*
  This self-propagating malicious code exploits a known Buffer Overflow vulnerability in IIS Indexing Service DLL. The victim hosts are scanned for TCP port 80 and the attacking host sends the exploit string to the victim, that if vulnerable and not yet affected is compromised by spawning threads to scan random IP addresses for hosts listening on TCP port 80, exploiting any other vulnerable hosts that can be found.

- *September 18, 2001 - Nimda (CA-2001-26)*
  The "W32/Nimda worm" or "Concept Virus (CV) v.5." appears to spread by multiple mechanisms: from client to client via email, from client to client via open network shares, from web server to client via browsing of compromised web sites, from client to web server via active scanning for and exploitation of various Microsoft IIS 4.0 / 5.0 directory traversal vulnerabilities.

- *January 25, 2003 - Slammer (CA-2003-04)*
  This worm, also known with the name "Sapphire", targeting SQL Server computers, is self-propagating malicious code that exploits a known vulnerability that allows for the execution of arbitrary code on the SQL Server 2000 computer due to a stack buffer overflow. Once the worm compromises a machine, it will try to propagate itself. The worm will craft packets of 376-bytes and send them to randomly chosen IP addresses on port 1434/udp. If the packet is sent to a vulnerable machine, this victim machine will become infected and will also begin to propagate.

- *August 11, 2003 - W32/Blaster (CA-2003-20)*
  The W32/Blaster worm exploits a vulnerability in Microsoft's DCOM RPC interface. Upon successful execution, the worm attempts to retrieve a copy of the file msblast.exe from the compromising host. Once this file is retrieved, the compromised system then runs it and begins scanning for other vulnerable systems to compromise in the same manner. In the course of propagation, a TCP session to port 135, 139 and 445 are used to execute the attack. The worm also includes the ability to launch a TCP SYN flood denial-of-service attack against windowsupdate.com.

- *Mar 19, 2004 - Witty (IN-2004-01)*
  The witty worm, targeting a buffer overflow vulnerability in several Internet Security Systems (ISS) products, was the first widely propagated Internet worm to carry a destructive payload. Once it infects a computer, it deletes a randomly chosen section of the hard drive, over time rendering the machine unusable. Witty was started in an organized manner with an order of magnitude more groundzero hosts than any previous worm and represents the shortest known interval between vulnerability disclosure and worm release – it began to spread the day after the ISS vulnerability was publicized.

- *Apr 30 - May 4, 2004 - Sasser - (VU#753212)*
  This worm attempts to exploit a buffer overflow vulnerability in the Windows Local Security Authority Service Server (LSASS). The vulnerability allows a remote attacker to execute arbitrary code with SYSTEM privileges. The worm has been reported to propagate by scanning random IP addresses on port 445/tcp to identify vulnerable systems. When a vulnerable system is found, the worm will exploit the LSASS vulnerability, create a remote shell on port 9996/tcp, and start an FTP server on port 5554/tcp. The victim system will then connect back to the attacking system on port 5554/tcp to retrieve a copy of the worm.

We also considered some other significant network infrastructure failure events affecting large geographic areas such as the World Trade Center attack in September 2001 and the North-East US and Italy blackouts in August and September 2003. We observed with a two hours or 15 minutes (only for BGP update messages) granularity for each of the above events the monthly trends of global routing table size, route damping and BGP update (announcement and withdraws) rates, to demonstrate a direct correlation between some of the worm spreading or power failure events and the global routing instability phenomena observed in the Internet during them. We included in our analysis only significant phenomena that show a duration of at least three hours and thus are clearly distinguishable from other localized infrastructure failure events not affecting global routing stability. In fact, it is common for individual peers to contribute with instantaneous "spikes" of high volumes of BGP advertisements on the IXP they are connected to, presumably reflecting BGP sessions closing and opening close to the collection point. During large-scale worm outbreak, instead, all peers experience a "wave" of smoothly increasing traffic and/or BGP activity, sustained for several hours. Our analysis has also been materially aided by data collected by the network security community and announced on several network security mailing lists and WWW resources, where we matched our observation with the available anecdotal reports of sudden connectivity losses, ARP storms, and similar localized worm effects.

4.2 The Most Interesting Results

The data resulting from our observations shows at a first glance a sustained and significant increase in the number of BGP damped routes registered on July 19^{th} 2001, correlated with the main propagation phase of the first large scale Windows-based worm, Code Red, which infected over 359,000 hosts on the Internet in less then a day. We can easily see this effect from Figure 1 below showing the number of damped routes sent and received by neighbor routers in the Route-Views and Ripe NCC monitoring points.

Such an impact on BGP seems to be surprising since the worm was directed against web servers while BGP

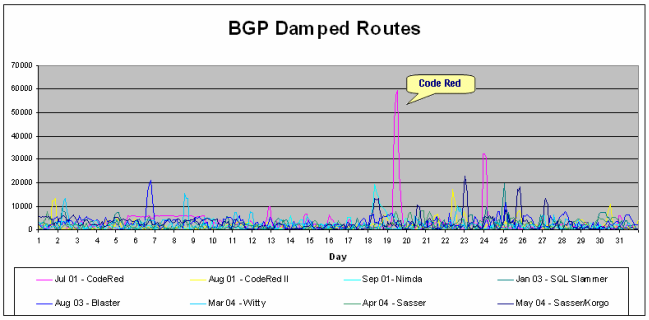

Figure 1: BGP damping rate observation

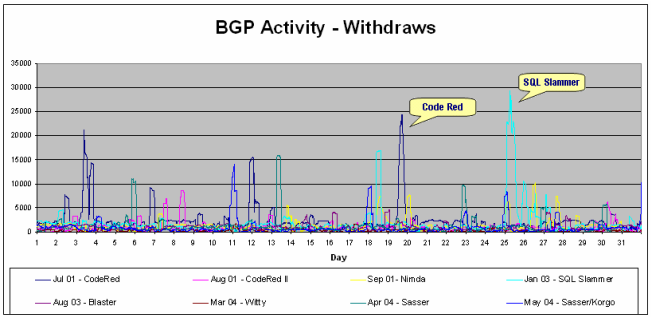

Figure 2: BGP withdraw activity

is running on routers which provide reachability to all networks. But, although this particular worm has not directly affected the routers themselves, as it infected a very large number of hosts on the network, simultaneously contributing to the worm propagation and thus generating a lot of traffic, it heavily impacted the routers indirectly through router failures due to excessive CPU and memory utilization, as well as through more obscure effects such as proactive reconfiguration or disconnection of certain routers by network administrators. All the above events generated continuous intermittent reachability failures on the affected destinations and hence route flapping that clearly correspond to the sustained spike in damping rate. It is interesting to see that the volume of worm traffic did not appear to be high enough to cause significant BGP session resets due to severe congestion. We also observed that the Code-Red worm attack was closely correlated in time with a large spike in the number of BGP route withdraw from multiple ISPs received at the observed monitoring points. Figure 2 below shows the sustained increase in the number of withdraw messages observed during the same worm propagation period. A withdraw message is sent when an AS no longer has any route to reach a particular prefix. The dramatic increase in the number of withdraws indicates that, not only did the paths to prefixes change during the worm period, but some prefixes were declared unreachable. This suggests that there was a loss of connectivity to some portions of the Internet.

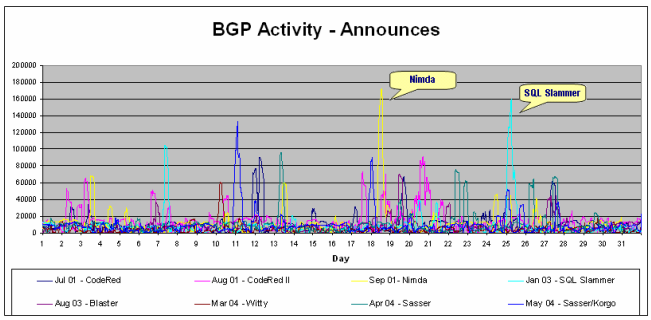

Figure 3: BGP announcement activity

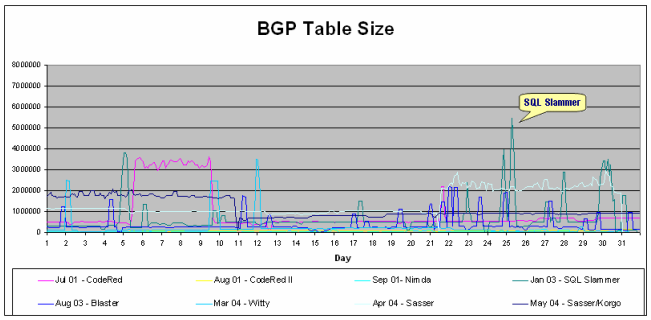

Figure 4: BGP routing table size

The rightmost and strongest event that can be observed from Figure 2, together with the Code Red, is related to the occurrence of the first Slammer outbreak, on January 25^{th} 2003. Slammer was the fastest computer worm in history and as it began spreading throughout the Internet, it doubled in size every 8.5 seconds. It infected perhaps 75,000 hosts, and more than 90 percent of them within the first 10 minutes, when it achieved its maximum scanning rate (over 55 million scans per second) [6]. Due to the congestion caused on the network transport backbone by its aggressive scanning activity and very fast propagation it taken out of service many destinations on the Internet, hence the large number of withdraws sustained for more that 24 hours. Combined and closely correlated with the increase in withdraw messages seen in Figure 2, we can observe in the following plot (Figure 3) a dramatic increase in the number of BGP route announcements received at the monitoring points. Early on January 25, the number of BGP updates spikes dramatically and tails off by January 26th.

Although large spikes of BGP updates are observed on a few other days as well, those on January 25^{th} 2003 and on September 18^{th} 2001 (together with the Nimda worm appearance) rose much higher and lasted longer. Such behavior was taken as an indication that the worm attack caused a significant global routing instability. Slammer spread nearly two orders of magnitude faster than Nimda, but it probably infected fewer machines. Both worms used the same basic strategy of scanning to find vulnerable machines and then transferring the exploitive payload; they differed in their scanning constraints. While Nimda was latency limited, Slammer was bandwidth-limited, allowing it to scan as fast as the compromised computer could transmit packets or the network could deliver them. Anyway, the scan activity of both the Nimda and Slammer worms might have contributed to intermittent congestion that led to frequent BGP session resets. In general, long-lasting, high-diversity, high-rate route churn can be produced by any mechanism that is causing a large number of BGP sessions to close and reopen repeatedly. Another similar effect also observed for Slammer worm when it was at most of its propagation is the some hours-long drop (transposed in absolute value we can see it as a spike) in the overall reachability rate (Internet routing table size) that can be observed from Figure 4.

Here too we can be observe the interesting effect of Slammer that with its extremely high scan rates and multiple attack modes generated very heavy traffic impacting the Internet through bandwidth consumption that reflects in the dramatic decrease of BGP available routes. The sustained and sharp drop in BGP routes at the first time of appearance of the Slammer worm, around January 25^{th} 2003, together with the previously presented evidences in BGP announcements and withdrawals, provides excellent correlation between the worm activity and Internet routing instability. Furthermore, Internet backbone disruption observed during the above outbreak was significant enough that various peering points became saturated due to worm traffic. Another important consideration is that, although worm attacks mainly affected connectivity at certain (or almost all) edges, whose intermittent reachability rippled through to the rest of the Internet as rapid BGP update exchanges, we obtained a clear evidence that, with the current BGP-based inter-domain routing design, a local change can indeed cause an unwanted global congestion effect. A truly resilient global routing system should not propagate local changes globally, in order to scale well.

Finally, it is worthwhile to note that the World Trade Center attack of September 2001 do not impact the collected data at all. The reason being that this event, while significant, did not impact overall Internet routing stability because it was a localized event and the network rapidly converged around the disaster scenario. The same considerations can be done for both the US Northeast and Italy blackouts that take place respectively in 14^{th} August and 28^{th} September 2003 where we only noticed a slight fall in global reachability (in term of number of routes in the full Internet routing tables) due to the large coverage of the geographic areas interested in the power failure. More generally, our examination of the data indicates that localized failures in the Internet core infrastructure (fiber cuts, flooding, power failures, building collapses) tend to generate only short-term increases in the external BGP prefix advertisement rate, which decline in a matter of minutes as the highly redundant peering in the Internet

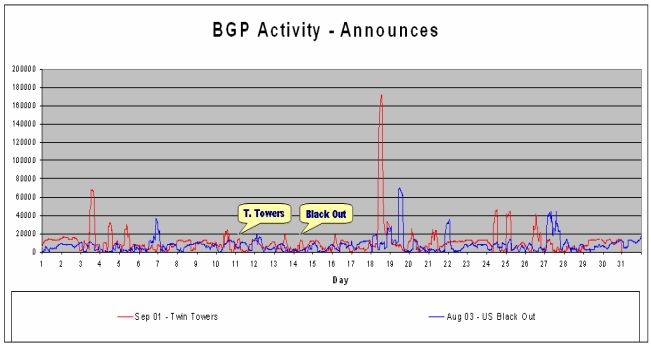

Figure 5: BGP activity during the WTC attack and US blackout

core topology routes around the damage, although specific networks may remain unreachable until the damage is repaired. The BGP announcement activity collected during September 2001 and August 2003, depicted in Figure 5 below. strongly evidences the above considerations.

4.3 The Dynamics Explained

The above events clearly demonstrated that an application layer event (the worm attack) caused problems at lower layers of the Internet infrastructure. Fortunately, only some of the observed worms triggered a widespread end-to-end routing instability that originated at the Internet edge, which refers to endpoint or stub networks attached to an ISP core and quickly affected the core itself. They were typically those that infected the greater number of hosts or were characterized by the most aggressive scanning behavior. The intense scanning activity of worms destroyed the locality of reference needed for the routers' caching system. Briefly, the worms started random IP port scanning and most of the random IP addresses generated were not in the cached entries of the routing tables of the involved routers, causing frequent cache misses, and in the case of invalid IP addresses, generation of ICMP host or network unreachable messages. Address Resolution Protocol (ARP) behavior was likewise affected, as local area networks received worm packets whose addresses were legitimate to subnet-work level, but whose full IP address failed to match any device in the subnetwork. All the above causes exacerbated the load already generated by the worm spreading traffic and rapidly led to router CPU overload, causing routing sessions reset or in the worst case router crashes. Router or routing sessions failure led to announcement withdrawal by the peers, generating a high level of advertisement traffic. When the router came back on, it required a full state update from it's peers, creating a large spike in the load of it's peers who provided the state dump. Once the restarted router obtained all the dumps, it dumped its full state to all its peers, creating another spike in the load. Frequent full state dumps led to more CPU overload, leading to more crashes, and the propagation of the cycle repeated until filters were applied to drop the attack traffic.

5 Related Work

The impact of worms or any kind of infrastructure stress events on global routing dynamics is still an almost unknown, or at least a scarcely understood, matter and to date there are very few measurement-based analyses [3, 4]. The impact of worms or any kind of infrastructure stress events on global routing dynamics is still an almost unknown, or at least a scarcely understood, matter and to date there are very few measurement-based analyses [5]. Their data came from core border routers, and they observed that an extremely large fraction of BGP updates were pathological. While we too have observed large volumes of BGP traffic, we taken more interest in anomalies in the collective dynamic behavior of BGP message traffic observed during some known stress-related events to demonstrate how and if they affected the overall network stability.

6 Conclusions

The impact of worms on the Internet has increased significantly over the past five years. In particular worms such as Code Red, Nimda and Slammer prove that the ability to effectively impact the overall Internet behavior by infecting hosts on a global scale in a matter of minutes is a reality. This impact is not only felt at the connection endpoint where the worm takes residence and replicates itself but also on the infrastructure in-between, and sometimes on the Internet core, due to the very aggressive spreading activity of the most recent worms on the available high speed links and to the more and more homogeneous installed software base that can be easily used for large scale exploitation. In the period of time that the most aggressive infections were at their most severe levels a unique effect on the overall connectivity began to be observed whereby global routing instability was detected throughout the Internet. Our analysis, however, revealed several weak points in the Internet routing infrastructure, such as BGP's sensitivity to the transport session reliability, its inability to avoid the global propagation of small local changes, and its certain implementation features whose otherwise benign effects only get amplified under stressful conditions. The focus in the future must be to build even greater resiliency and adaptive containment countermeasures into the Internet infrastructure operating real-time smart checks on the network stability to prevent such events from recurring with even more impact.

References

[1] E. Cooke, Z. M. Mao, and F. Jahanian, *Worm Hot-spots: Explaining Non-uniformity in Worm Tar-*

geting Behavior, Technical Report CSE-TR-503-04, University of Michigan, 2004.
[2] CERT Coordination Center: http://www.cert.org
[3] J. Cowie and A. T. Ogielski, *Global Routing Instabilities Triggered by Code Red II and Nimda Worm Attacks*, Renesys Corp., Technical Report, Dec. 2001.
[4] C. Labovitz, G. R. Malan, and F. Jahanian, "Origins of Internet routing instability," in *Proceedings of INFOCOM 1999*, pp. 218-226, 1999.
[5] C. Labovitz, G. R. Malan, and F. Jahanian, "Internet routing instability", in *Proceedings of ACM SIGCOMM 1997*, pp. 115-126, 1997.
[6] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Stani-ford, and N. Weaver, "Inside the slammer worm," *IEEE Security and Privacy Magazine*, vol. 1, no. 4, pp. 33-39, July 2003.
[7] C. Nachenberg, "The evolving virus threat," in *23rd NISSC Proceedings*, Baltimore, Maryland, 2000.
[8] Y. Rekhter and T. Li, *Border Gateway Protocol 4*, IETF RFC 1771, July 1995.
[9] RIPE RIS Project: http://www.ripe.net/projects/ris/
[10] J. Rochlis and M. Eichin, "With microscope and tweezers: The worm from MIT's perspective," *Communications of the ACM*, vol. 32, no. 6, pp. 689-698, June 1989.
[11] E. Spafford, "The Internet worm: crisis and aftermath," *Communications of the ACM*, vol. 32, no. 6, pp. 678-687, June 1989.
[12] University of Oregon Route Views Project: http://www.routeviews.org/
[13] US-CERT: http://www.us-cert.gov
[14] C. Villamizar, R. Chandra, and R. Govindan, *BGP Route Flap Damping*, IETF RFC 2439, Nov. 1998.

Francesco Palmieri holds two Computer Science degrees from Salerno University, Italy. Since 1989, he worked for several international companies on a variety of networking-related projects, concerned with nation-wide communication systems, network management, transport protocols, and IP networking. Since 1997 he leads the network management/operation centre of the Federico II University, in Napoli, Italy. He has been closely involved with the development of the Internet in Italy in the last years, particularly within the academic and research sector, as a member of the Technical Scientific Committee and of the Computer Emergency Response Team of the Italian Research Network GARR. He is an active researcher in the fields of high performance/evolutionary networking and network security. He has published several papers in leading technical journals and conferences and given invited talks and keynote speeches.