SANDIA REPORT

SAND2003-1866
Unlimited Release
Printed June 2003

Distributed Denial-of-Service Characterization

Michael Berg, Philip Campbell, Timothy Draelos, David Duggan, Mark Torgerson, Brian Van Leeuwen, and William Young

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94AL85000.

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865)576-8401
Facsimile: (865)576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.doe.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800)553-6847
Facsimile: (703)605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2003-1866
Unlimited Release
Printed June 2003

Distributed Denial-of-Service Characterization

Timothy Draelos and Mark Torgerson
Cryptography and Information Systems Surety Department

Michael Berg, Philip Campbell, David Duggan, Brian Van Leeuwen, and William Young
Networked Systems Survivability and Assurance Department

Sandia National Laboratories
P. O. Box 5800
Albuquerque, NM 87185-0785
{tjdrael, mdtorge, mjberg, plcampb, dduggan, bpvanle, wfyoung}@sandia.gov

Abstract

Distributed denial of service (DoS) attacks on cyber-resources are complex problems that are difficult to completely define, characterize, and mitigate. We recognize the process-nature of DoS attacks and view them from multiple perspectives. Identification of opportunities for mitigation and further research may result from this attempt to characterize the DoS problem space. We examine DoS attacks from the point of view of 1) a high-level that establishes common terminology and a framework for discussing the DoS process, 2) layers of the communication stack, from attack origination to the victim of the attack, 3) specific network and computer elements, and 4) attack manifestations. We also examine DoS issues associated with wireless communications. Using this collection of views, one begins to see the DoS problem in a holistic way that may lead to improved understanding, new mitigation strategies, and fruitful research.

Contents

Figures

1 Introduction

The problem of cyber-attacks has many people and organizations anxious about the potential damage of future attacks and the safety of conducting business on the Internet. We know that adversaries continue to develop and explore innovative ways to enhance the effect of their attacks. One way of multiplying the effect of an attack is by taking advantage of the distributed nature of the Internet. The class of attacks known as Denial-of-Service (DoS[1]) can benefit from a multiplicative effect, resulting in network or distributed Denial-of-Service (DDoS[2]) attacks. Therefore, DDoS can be considered a special case or subset of DoS attacks. For the remainder of the paper, we use the DoS acronym to include DDoS.

Availability attacks (i.e., DoS attacks) are fundamentally different than other cyber-attacks that involve the **presence** of a listener or a "bit-flipper." DoS attacks involve the **absence** of a capability, of data, or of a service. Perhaps it is this contrarian nature that makes the problem difficult to fully understand and solve.

The central goal of this paper is an attempt to provide insight into characterizing DoS attacks through a collection of different perspectives. The primary aspect of DoS attacks that motivates our approach to DoS characterization is the following.

> *DoS attacks involve a process, not a single event, that comprises multiple activities through time and space from its origin to its victim.*

Our goal is to holistically characterize the entire DoS attack process, not just the endpoint of attacks, from multiple viewpoints, not unlike the way physicians examine the same human body with different instrumentation, such as X-ray, ultrasound, and magnetic resonance imaging (MRI) equipment. Various views can be used to examine the DoS problem space, to identify areas in need of safeguards, and potentially to enable one to predict future DoS activities.

Taken together, our collection of views offers a valuable way of characterizing the DoS problem space. The set of perspectives serves as a toolbox for those interested in DoS prevention. The intent is to characterize the DoS problem space such that current research efforts can be evaluated in terms of their relevance to DoS mitigation and new, fruitful research areas can be identified.

The discussion thus far has not distinguished between the effects of DoS on wired and wireless networks. Although our approach is pertinent to both forms of communication, we introduce wireless DoS issues here and present details in Section 6. When considering how DoS attacks can be used against wireless networks, several key factors differentiate wireless networks from their wired counterpart. The primary factors are the use of free space as the communication channel in a wireless network, typically RF, optical, or infrared, and the mobility and resource restrictions of many wireless devices. The frequency spectrum of the wireless channel is a scarce resource that a DoS attack can consume. Mobile wireless nodes present routing challenges in an ever-changing network topology, and power restrictions limits the available protocols and mitigation solutions. For these reasons, wireless networks are an attractive target for DoS attacks, and a natural area for DoS attacks to migrate to in the future.

---

[1] A simple definition of Denial of Service is an attack designed to render a computer or network incapable of providing normal services [20].

[2] A Distributed Denial of Service attack uses multiple computers to launch a coordinated DoS attack against one or more targets [20].

The remainder of the paper presents the DoS process from various views used to characterize the DoS problem space. We present four different views of DoS, summarized in Table 1-1, which references the section in which they are presented in the paper.

View	Viewpoint	Section
High-Level	A high-level representation of the network architecture and information flow of DoS processes.	2
Communication Stack	A view that tracks attack activity through the communications stacks from origination to the victim. With this view, one can clearly see the problematic layers of the communications stack that deserve attention.	3
Network and Computer Elements	A view that identifies tangible elements of a network or computer that an adversary attacks. Elements in particular need of safeguards are identified with this view.	4
Attack Manifestation	A view that distinguishes between attacks by their DoS manifestations.	5

Table 1-1 Collection of Views of DoS attacks.

In addition to the multiple views of DoS attacks, the paper addresses DoS issues associated with wireless networks or network elements in Section 6, presenting specific points of concern at the physical, data link, and network layers of communication. In Appendix A, we provide an annotated review of a significant body of work related to DoS. Appendix B includes a mapping of DoS attacks to the Communication Stack View and Appendix C contains a list of DoS attacks with brief descriptions.

2 High-Level View

The High-Level View provides common terminology and a general network framework for describing the DoS attack process. It is used as a basis to help explain the other views. This view attempts to offer a means to represent all DoS attacks from their origin to their intended victim. However, we recognize that it is unlikely that one will be able to partition a network into disjoint, descriptive, and useful categories that *a priori* cover every conceivable attack scenario.

One of our goals is to provide a broader understanding of possible DoS attacks. To help with this, we loosely define three networks that correlate with the nature of the attacks. One may view a DoS attack as a process that has a beginning, an end goal, and which may or may not utilize resources acquired during the process of the attack. With this in mind, we divide the participating nodes into three sub-networks (see Figure 2-1). These three networks are not a separation based on physical location, but rather a logical description of their roles in the attack.

1. **Origination Network** – The Origination Network is the collection of one or more nodes that the adversary has access to and uses to initiate the DoS attack. The adversary is not limited to a single individual, and in the cyber realm, physical control of a node is not a requirement for access to the node.
2. **Facilitation Network** – The Facilitation Network is the collection of zero or more nodes used by the adversary after an attack's initiation to help accomplish the attack on the Victim Network.
3. **Victim Network** – The Victim Network is the collection of one or more nodes to which the adversary wishes to deny service.

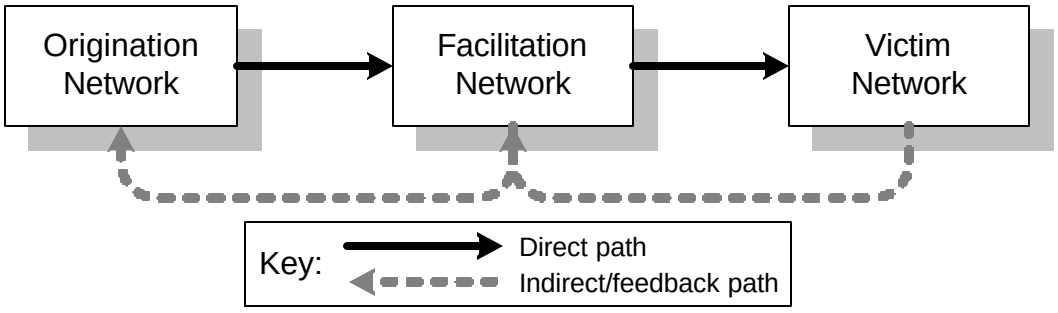

Figure 2-1 The process flow and logical architecture of DoS attacks.

The black arrows in Figure 2-1 emphasize the primary propagation of attacks from the Origination Network to the Victim Network, while the gray arrows indicate that attacks can feed back and manifest themselves in many nodes during various stages of an attack[3].

Note there may be some debate as to when a particular attack is originated. For instance, in January the adversary uses their computer to plant code on another computer in the Facilitation or Victim Network. In July the planted code wakes up and begins its adversarial action. From one perspective, it may appear that the attack was initiated in July from the compromised node. On the other hand, one may view the attack as being initiated from the adversary's computer in January. We take the latter point of view.

Because of the progressive nature of an attack, the role that any network node plays in the attack may change as the attack proceeds. The adversary's node in the Origination Network may be directly connected to the node he wishes to attack. In this case, there would be no Facilitation Network. On the other hand, if the attack were forwarded to the Victim Network by an unwitting router, then the unwitting router would be a member of the Facilitation Network. Often the Facilitation Network is an intermediate target for the attacker, used to distribute or multiply the DoS effect on the Victim Network.

During the course of an attack, many nodes may be compromised and/or harmed in various ways, but the adversary will have some final goal or primary victim in mind. It is the final goal that distinguishes the Victim Network from the intermediate or Facilitation Network. Even if the goal is to foment chaos in the larger network, one can appropriately identify a Victim Network.

It is important to note that for any given attack, there will be an Origination Network and a Victim Network, but there may not be a Facilitation Network. Any particular node may be in any one or more of the three networks as an attack progresses. One benefit of viewing the three networks

---

[3] As an example, many of the flooding attacks use spoofed source IP addresses, often of non-existent hosts. Reply traffic from the Victim Network will eventually end up as "undeliverable" somewhere in the Facilitation Network. Depending on the type of packets used in the attack and on router and firewall configurations along the way, there may also be ICMP "host unreachable" messages sent from a router in response to the undeliverable reply traffic coming from the Victim Network. Another example is that some DoS attacks require the attacker to hold an open connection or to interact with a daemon (like sending a large block of malformed input to crash or lock it). In these cases there will be communication with the Victim Network with response traffic going back through the Facilitation Network and to the Origination Network.

from a logical rather than a physical point of view is that we do not need to make a distinction between nodes that occupy more than one sub-network category and those that do not.

The High-Level View provides the following potential benefits.

1. It presents an inclusive framework for DoS attacks to which more detailed views can relate.
2. It depicts the process nature of DoS attacks. As time advances during a DoS attack, different elements of a network become active. For example, the attack may manifest itself in the Facilitation Network, then the Victim Network, and again in the Facilitation Network (i.e., an attack is a time sequence of events).
3. It allows the coverage of other views to be identified. The three views presented in Sections 3 through 5 provide broad coverage of the DoS network space (Origination, Facilitation, and Victim), as seen in Figure 2-2.

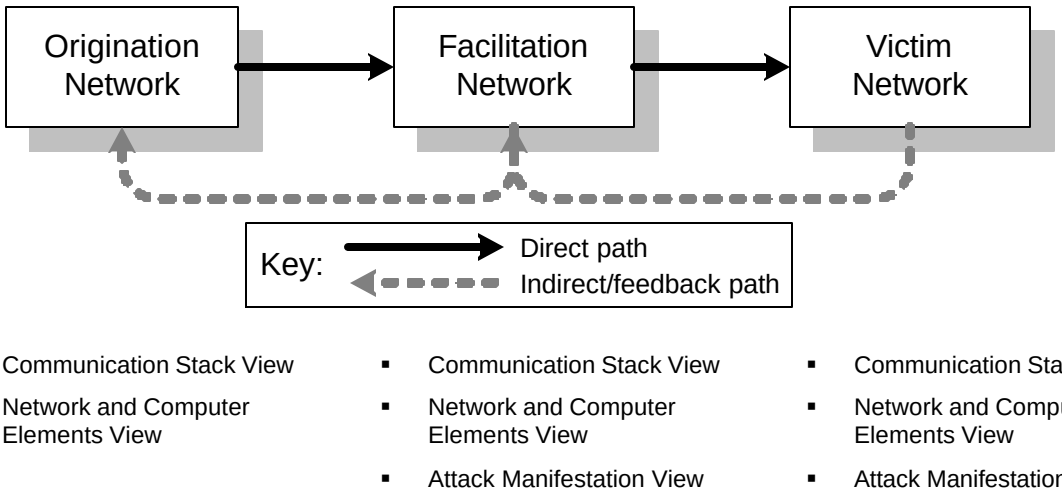

Figure 2-2 Coverage of the DoS network space by our characterization views.

3 Communication Stack View

The basic concept behind the Communication Stack View is to describe DoS attacks by behavior exhibited at the various layers in a communication stack reference model as the attack progresses from the Origination Network, through the Facilitation Network, to the Victim Network. This perspective helps identify where potential mitigation strategies must exist within the communication process in order to defend against or detect a particular attack or set of attacks. In addition, by developing this view of attacks, classes of attacks with similar behavior may lead to a natural classification of attacks in general. A general classification of attacks supports the development of mitigation strategies aimed at protecting against more than a single attack.

Several communication stack models are available to support the desired mapping. The two most common are the Open Systems Interconnection (OSI) and the TCP/IP (sometimes called the Internet) Models. The OSI Model includes two layers not often distinguished in a communication process, i.e. the Presentation and Session Layers, while the TCP/IP Model only utilizes four levels of granularity. Although either of these reference models is adequate, the hybrid model, found in [21], provides a useful compromise of the two previous models (see Table 3-1). Our communication stack DoS characterization view utilizes the hybrid model.

Application Layer
Transport Layer
Network Layer
Data Link Layer
Physical Layer

Table 3-1 Hybrid OSI/TCP/IP Layer Model**.**

The goal is to identify which layers are utilized in the DoS attack, and whether the layers exhibit normal or abnormal behavior. A critical element of this approach is the distinction between normal and abnormal behavior. In this current discussion, abnormal behavior refers to the use of a protocol or process in a manner not intended in its original design. For example, using TCP to establish a series of half-open connections constitutes abnormal behavior of the protocol. It should be noted that some attacks will not exhibit an abnormal behavior at any layer in the stack, but determining the layers involved can still provide benefits. Some initial questions required to map an attack to the communication stack model include the following.

1. At what layer does the attack originate?
2. At what layer is the attack aimed at on the Victim Network?
3. Which layers in the Facilitation Network require either normal or abnormal behavior?

The "Hybrid Model" is replicated for each network of the High Level View, described in Section 2. Then, given an individual DoS attack, the communication stack behavior is mapped for each of these networks. The mapping depicts normal or abnormal behavior at the various layers of the communication stack. Note that some layers may not be utilized in a given attack, and will be left blank in the mapping process. An example of this approach is depicted in Figure 3-1. Table B-1 in Appendix B presents the same information in a compact manner for a list of known attacks, and provides a more effective display of the mappings.

Mapping DoS attacks to this hybrid reference model provides the following potential benefits.

1. It allows identification of the potential coverage by a proposed solution. For example, a mitigation approach at a layer higher than the attack will likely not solve the problem.
2. It helps in identifying the coverage or scope of current attacks, i.e. percentage exploiting the network layer versus the transport layer.
3. It supports the general classification of DoS attacks.
4. It may provide insight into which attacks targeted for a wired network can easily transition to a wireless network.
5. It provides a graphical representation of DoS attacks.

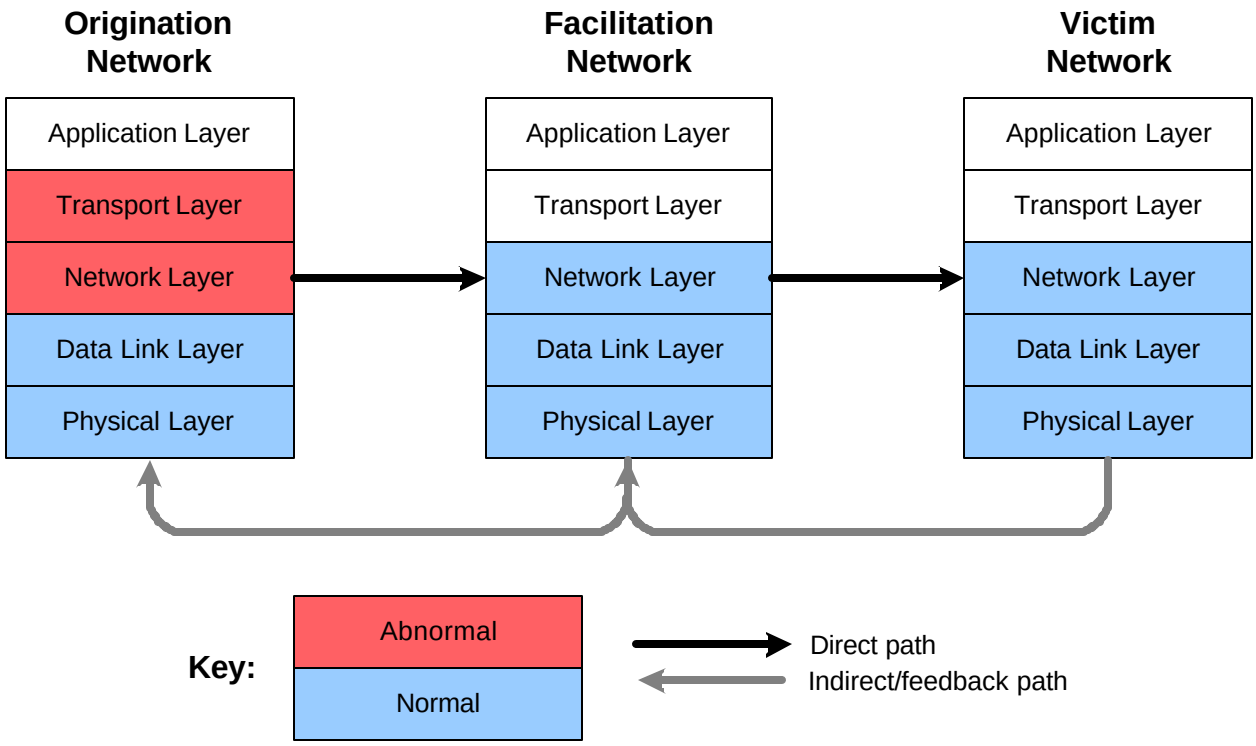

Figure 3-1 The Communication Stack View of a TCP/SYN Flood attack.

4 Network and Computer Elements View

While the Communication Stack View identifies attack behaviors in the abstraction of communication layers, the Network and Computer Elements View identifies the specific elements of software, hardware, or protocols where the attacks take place. The Network and Computer Elements View of DoS characterization focuses attention on specific and tangible, as opposed to logical, elements of a system, protocol, or data packet that allow the adversary to achieve their goal, namely denial of some service.

In this section, the word target is used to describe an entity which an attacker plans to compromise. This entity may not be the service that the attacker wishes to deny, but is the element of the network which will allow their ultimate purpose to succeed. In other words, the target does not equate to the Victim Network. For example, using a military illustration, targets that will deny fueling service to automobiles include refineries, gas stations, and gas trucks, but only one of these can be the target of a given missile. Therefore, the target is where the "explosion" or compromise occurs. Given that a DoS attack can consist of multiple exploits, there can exist multiple targets spread out over time and space. With this view of the attack space, we recognize that adversaries identify and attack real elements of networks and computers for compromise, whether hardware, software, or firmware.

Howard and Longstaff [7] include the following elements in their list of targets.

1. **Account –** a domain of user access on a computer or network, which is controlled according to the user's account name, password, and use restrictions.
2. **Process –** a program in execution, including data, and program registers.
3. **Data –** representations of facts, concepts, or instructions.

4. **Component –** one of the parts that make up a computer network.
5. **Computer –** a device consisting of processing units and peripheral units, controlled by internally stored programs.
6. **Network –** an interconnected or interrelated group of host computers, switching elements, and interconnecting branches.
7. **Internetwork –** a network of networks.

Although the first three target elements in the list above are logical entities and the rest are physical entities, they are all realized on physical components. This fact makes them targets for exploitation during an attack. The adversary will likely aim an attack at specific elements of a network with known vulnerabilities, whether in the Facilitation or Victim Network.

The above list of targets is generally not detailed enough to be helpful in discerning solutions to the DoS problem. Therefore, we use the following items as an example of the level of detail necessary to enable identification of likely targets and therefore areas where safeguards are needed.

1. **Communication Protocol –** Methods used for inter-computer communication whether implemented in software or hardware.
2. **Data –** Representations of facts, concepts, or instructions, such as configuration information, routing tables, and cryptovariables.
3. **Application –** Software with which users interact that is designed to carry out specific tasks, such as a web browser or mail server/client.

Figure 4-1 presents the Network and Computer Elements View diagrammatically, providing details of specific targets. The names of example DoS attacks are placed beside the network or computer element that is compromised during the attack (see Appendix C for a list of DoS attacks). Some attacks (e.g., TCP Syn Flood and UDP Flood) show up multiple times in the diagram because they assault multiple elements during the attack.

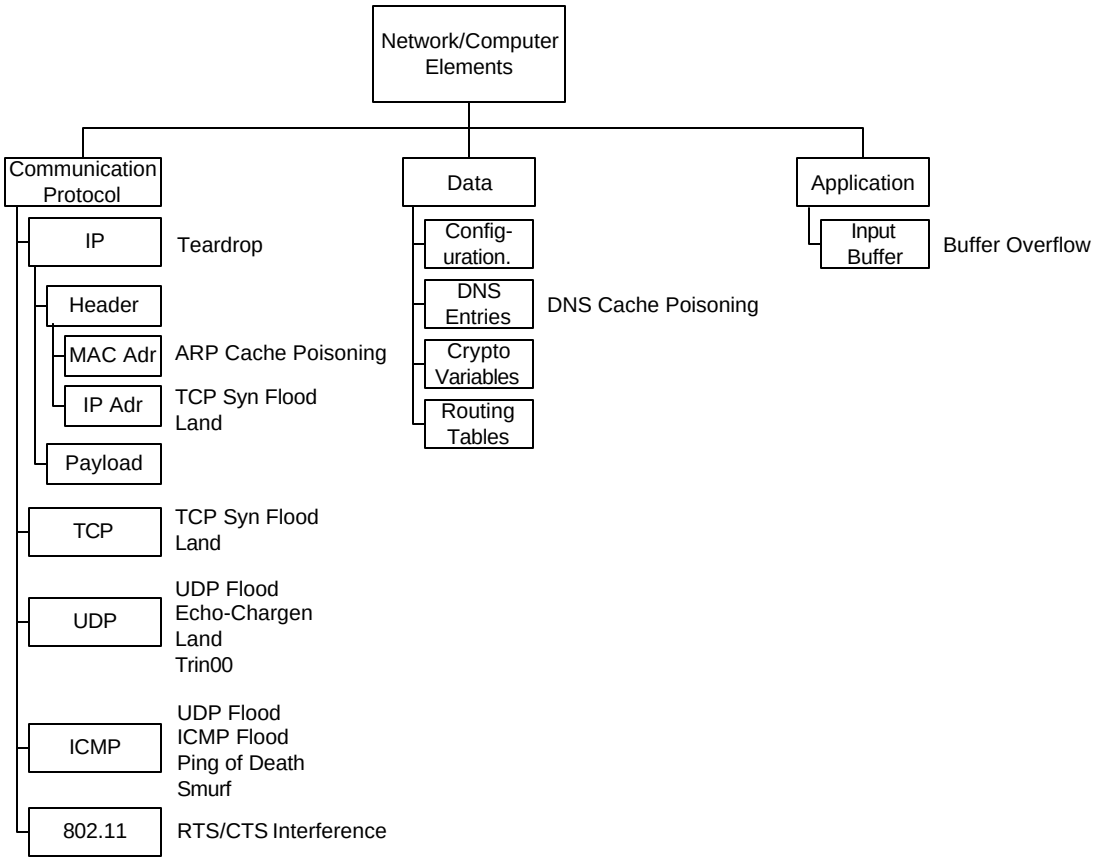

Figure 4-1 Network and Computer Elements View of specific DoS attacks.

Many of the identified network and computer elements will exist on machines in the Origination, Facilitation, and Victim Networks. However, given that a DoS attack is process-oriented, we observe the following.

1. DoS attacks aim their assault at different network and computer elements at different times during the attack.
2. A DoS attack may utilize network and computer elements in different ways, for example when intermediate targets are used to deny service on the Victim Network. Attack manifestations will be different, depending on the specific target and its role in the attack.
3. As the Internet changes and components are added, targets and methods will likely change.

Characterizing DoS attacks by network and computer element targets provides the following potential benefits.

1. It promotes consideration of specific safeguards to protect specific elements in a network, similar to the way we protect physical property. For example, people protect individual elements of their houses (e.g., doors and windows) in specific ways, such as locks, deadbolts, break sensors, etc.
2. It helps identify popular targets of current attacks as well as potential targets that aren't currently attacked.

While this view of DoS characterization identifies particular elements of a computer/network in need of protection, we also recognize that other, more process-oriented, systemic safeguards are important for raising the security of the Internet against DoS attacks. An example of a process-oriented safeguard is the use of intrinsically secure programming languages that don't allow buffer overflow programming errors. Other important process-oriented safeguards are enforcement of proper security procedures and installation of available security patches. These practices would raise the security level throughout the Internet, making it more difficult for DoS and other attacks to succeed.

5 Attack Manifestation View

While the Communication Stack View and Network and Computer Elements View identify **where** attacks occur, the Attack Manifestation View attempts to identify **what** occurs as a result of the attack. The view offered here categorizes DoS attacks according to the manifestation mechanism utilized. DoS attacks can manifest themselves in different ways in the Origination, Facilitation, and Victim Networks. Our focus is on the manifestations evident in the Victim Network as a result of a successful DoS attack.

By understanding the various manifestations of DoS attacks, work can be done to develop mitigations for those specific areas or to direct research to critical areas not being investigated. Mitigations can potentially be designed for entire classes of DoS attacks irrespective of the delivery mechanism used during the attack.

We present three diagrams that provide visuals for this view. Appendix C contains a list of selected DoS attacks with a description of how each attack relates to this view.

Top-Level Diagram

The highest-level diagram, shown in Figure 5-1, defines two branches that cover the manifestation mechanisms of DoS attacks. They are Resource Exhaustion, shown in Figure 5-2, and Service "Gone", shown in Figure 5-3. The Resource Exhaustion branch is based on the perspective of the service being provided, whereas the Service "Gone" branch is based on the perspective of the user of the service.

Resource Exhaustion Diagram

Under the Resource Exhaustion branch, all sub-branches deal with manifestations that occur by some resource being exhausted. The diagram in Figure 5-2 shows different resources that can be exhausted. This is an initial list and includes many of the resources that have been exhausted in previously identified DoS attacks. Additional resources we have identified could be useful to an adversary in the future.

Service "Gone" Diagram

The Service "Gone" diagram (Figure 5-3) shows the branches where the service *appears* to be gone. There are many different ways to make the service disappear. Listed in this branch are those that have been used in the past, as well as those identified through security studies that could possibly be used in the future. Making the service appear to be gone does not necessarily require the manifestation to be seen at the physical location of the service provided (e.g., the Victim Network). It could be such that the attack logically isolates the service from users of the service through other than direct means (e.g., via the Facilitation Network).

Characterizing DoS attacks by manifestation mechanisms provides the following potential benefits.

1. It promotes the detection of DoS attacks based on the evidence from manifestations. This could be helpful for early warning of attacks and for automated response to attacks.
2. It helps in identifying safeguards that can prevent an attack or class of attacks.
3. It identifies possible future attacks that exhaust resources or make services unavailable.

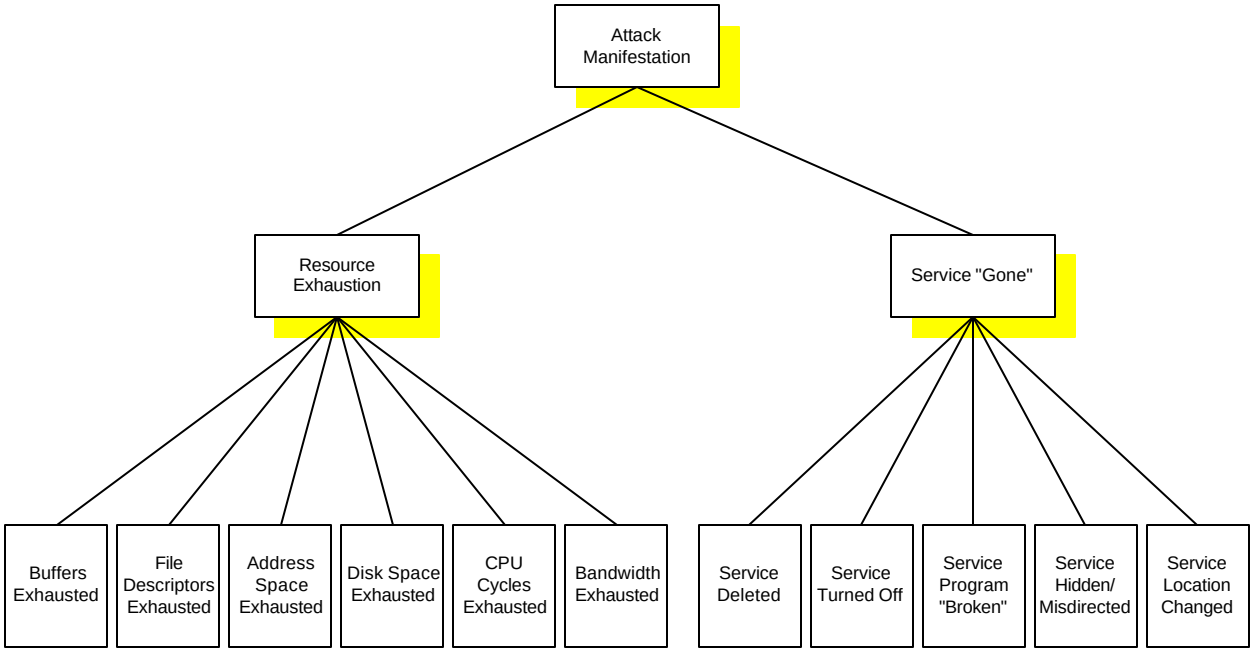

Figure 5-1 Attack manifestations.

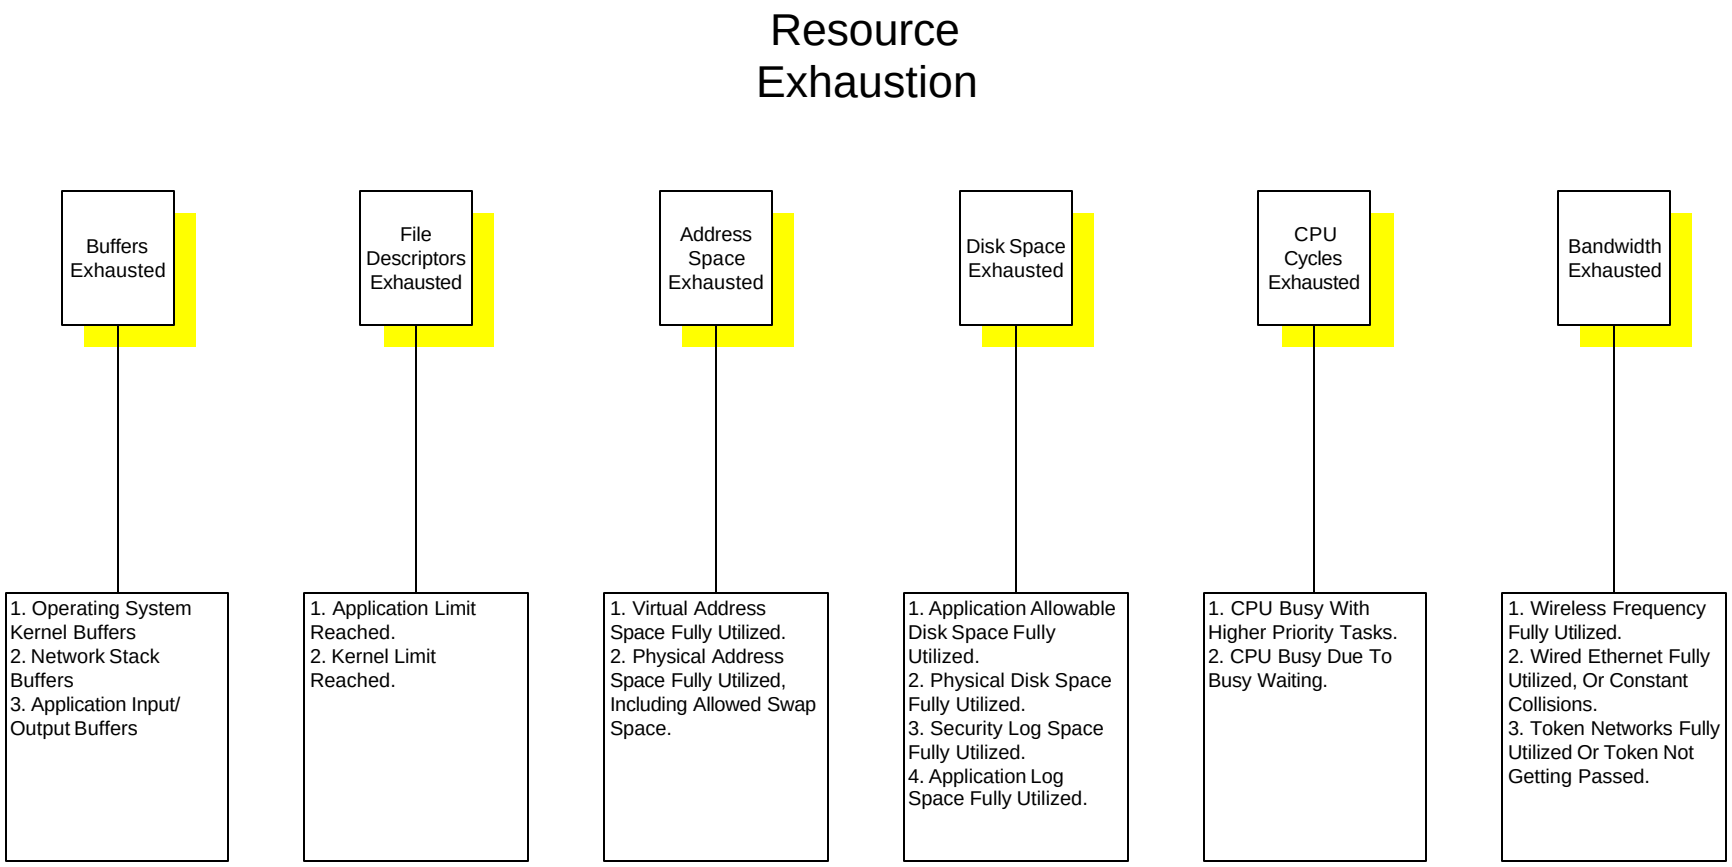

Figure 5-2 Resource Exhaustion manifestations.

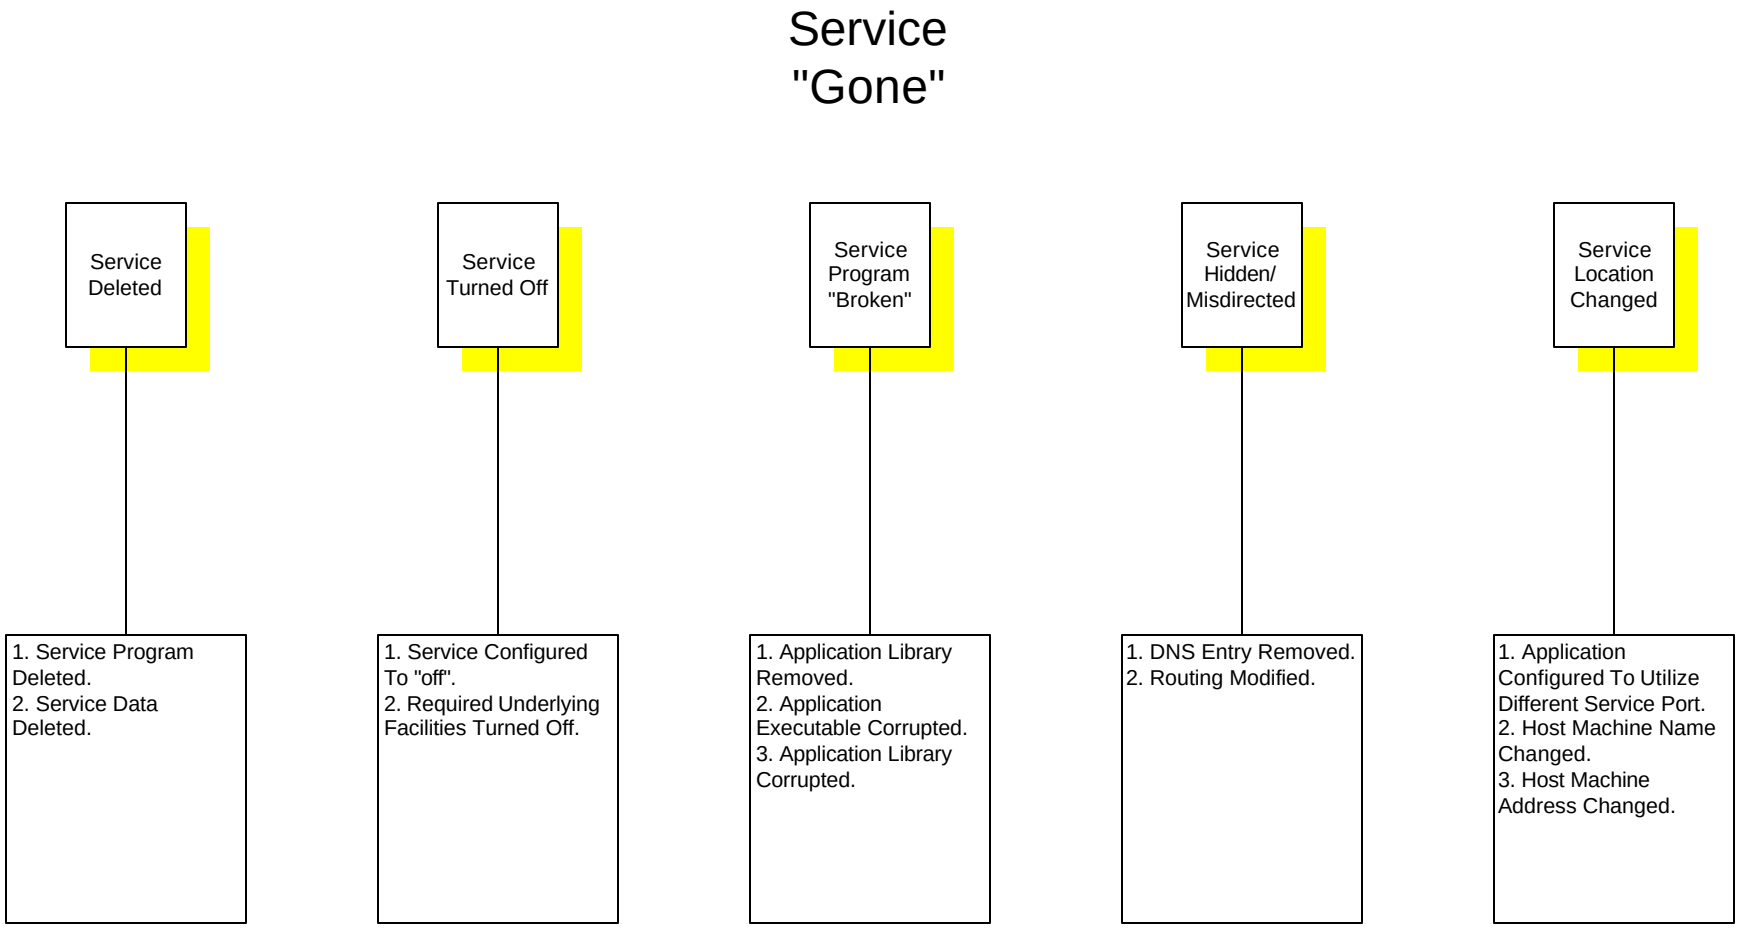

Figure 5-3 Service “Gone” manifestations.

6 Wireless DoS Issues

In this section, we change course from presenting different views that characterize DoS attacks to discuss DoS issues specific to wireless networks. While the views presented in the previous sections are applicable to both wired and wireless networks, the issues and potential DoS threats associated with wireless communication deserve special attention. The following discussion is presented on the three lower layers of the network stack from an adversary's point of view.

While the number of reported wireless DoS attacks to date is limited, such attacks are starting to occur. For example, the *SMS Flooder* [19] attack attempts to send huge volumes of SMS (Short Message Service) messages to a selected mobile device. As the use of wireless connectivity increases, we anticipate an increase in wireless specific, as well as hybrid (i.e. networks with both wired and wireless communication components) DoS attacks. While attacks will occur on both wired and wireless networks, differentiating characteristics between these networks affect the adversary's choice of attack, as well as the potential impact. In the next several paragraphs, we elaborate on some key characteristics of wireless networks that make such networks attractive to DoS attacks.

Wireless Characteristics

The primary factor that distinguishes a wireless network is its use of free space as the communication channel. Wireless networks use the radio, optical, or infrared frequency spectrum to transmit data across the wireless channel. A DoS attack can consume the frequency spectrum of the wireless channel without the need to break into the victim's system.

Mobile and self-contained nodes with severe resource constraints provide another key differentiating factor often associated with wireless networks. Distinguishing characteristics of a wireless mobile node (as compared to a typical wired node), include reduced processing capability, limited memory resources, battery provided power, and restricted bandwidth for communications. The resource limitations impact mobile nodes in two important ways.

1. Scarce resources can be depleted more quickly by DoS attacks.
2. Scarce resources limit the mitigation strategies that can be implemented.

Since DoS attacks typically focus on resource consumption, we expect that mobile wireless networks are especially vulnerable to these attacks. In the setting of the High-Level View, (Section 2), this may mean a DoS attack against a mobile node requires fewer Facilitation Network nodes. An attacker could also employ wired Facilitation Networks to attack the wireless portion of a heterogeneous network. In this case, the attacking network is not limited by the same resource constraints as the Victim Network, and could easily consume the resources of a much larger number of mobile, wireless nodes.

In the remainder of this section, we point out particular problems and some potential exploits attributed to the Physical, Data-Link, and Network layers of wireless communications. Attributes at those layers further emphasize why solutions applicable to the wired DoS problem will likely not directly transfer to the wireless environment.

Physical Layer

a. The wireless channel is an open medium that suffers from interference by competing users, natural interference, or malicious RF energy into the channel.
b. Typically, there is no physical barrier protecting access to the channel.

c. In most cases, the wireless channel can be accessed with the use of appropriate hardware and a gain antenna.
d. Non-invasive access to the wireless channel makes physical layer attacks more difficult to prevent in a wireless network.

Data-Link Layer

a. The data-link protocols control node access to the channel. Potential exploits include malicious nodes keeping the wireless channel busy by transmitting bogus channel access messages or directing spurious data for relay to a particular node to consume battery power.
b. The IEEE 802.11 standard represents a common wireless medium access approach, and uses both Collision Sense Multiple Access (CSMA) and Distributed Coordination Function (DCF) to control media access. The DCF is a virtual carrier sense that solves the hidden and exposed terminal problem[4]. DCF uses Request-to-Send (RTS), Clear-to-Send (CTS), data, and acknowledgment packet exchange to access the wireless channel and transmit data. Any disruption to the messages in this exchange can lead to DoS. For example, a node could send a continuous stream of RTS messages, forcing nodes to respond with CTS messages.

Network Layer

a. Network layer attacks against wireless network mobile nodes participating in message routing will impact the ability of the network to forward messages in a dynamically changing network. Maintaining correct routing tables requires the exchange of control messages among the network nodes, and disruption of these routing tables can cause a DoS.
b. Wireless LANs provide access to wired networks through a dual-membership host, namely a wireless access point. Compromise of the wireless access point can disrupt the communication ability of connected wireless nodes, both between the wired network and among the nodes (if the access point provides coordination functions).
c. Record and replay of control messages by cooperating malicious nodes in distant regions can create a spoofing condition indicating false neighboring nodes [22][18]. In this exploit, the adversary provides an out-of-band channel for data transfer between "neighboring" non-malicious nodes. When the adversary drops the out-of-band channel, the network will partition if it spreads beyond communication range.

In conclusion, wireless networks are susceptible to many of the same exploits that impact wired networks. However, in wireless networks the attacks may be easier to initiate since the wireless channel is easier to access. In addition, wireless devices are often equipped with unique protocols, such as Bluetooth™, SMS, IEEE 802.11, etc. As wired and wireless networks become tightly integrated, the interaction of these protocols with other wired network protocols might introduce new vulnerabilities. Application of the views presented earlier in this paper to wireless and hybrid networks will help identify key characteristics of potential wireless DoS attacks.

---

[4] In a wireless network that uses a common medium, each transceiver may not receive from, or detect the presence of, other transceivers. Thus carrier sensing alone cannot adequately regulate media access to prevent collisions at the receivers. The *hidden terminal problem* occurs when the medium is free near the transmitter and not free near the intended receiver, thus causing a message collision at the receiver. The *exposed terminal problem* occurs when the medium is busy near the transmitter and is free near the intended receiver, resulting in a lost opportunity to send a message.

7 Conclusions

With every new service offered via the Internet, the opportunities for its denial abound. Our approach to addressing this challenging DoS problem has been to examine attacks holistically, from beginning to end and from different points of view. Judging from the literature, a single view has not been able to solve the DoS puzzle. With a multiplicity of views, our goal has been to provide new insights and additional understanding of DoS attacks. This multiview perspective allows one to see how a suite of potential safeguards at various locations and levels of internetworking might best defend against DoS attacks. With this approach, diagnoses of DoS attacks may be possible based on observed network symptoms similar to the way a medical index is used to diagnose diseases.

We offer a way of looking at DoS attacks that recognizes DoS as a multistep, multiple-resource process. So while opportunities may abound for DoS attacks, opportunities for mitigation and detection also exist throughout the DoS process. We present various views of examining the DoS problem space.

1. High-Level
2. Communication Stack
3. Network and Computer Elements
4. Attack Space Manifestations.

Each of these views offers a different perspective on the DoS problem and identifies areas of concern where safeguards should be employed.

In general, wireless networks are not substantially different than wired with respect to DoS, so the findings in this paper are appropriate for both communication media. Although an abundance of wireless DoS attacks do not currently exist, the environment is rich with opportunity for just such attacks, and we expect a greater impact of DoS on wireless networks in the near future.

The following list of future work will help extend the concepts and conclusions presented here.

1. Further explore the best use of each view. This is an exercise in depth and may involve adding additional detail as well as suggesting mitigation solutions and predicting new attacks.
2. Explore additional views. This is an exercise in breadth and depth of each additional view, expanding the amount and quality of information for diagnoses of DoS attacks.
3. Map specific attacks to each view in such a way that a historical trend of DoS attacks may be discovered. This may lead to the prediction of new attacks.
4. Explore the interworkings of the suite of views. The goal would be to achieve an output that is greater than the sum of the parts.

8 References

[1] W. J. Blackert, D. C. Furnanage, Y. A. Koukoulas. “Analysis of Denial of Service Using An Address Resolution Protocol Attack.” Proc. of the 2002 IEEE Workshop on Information Assurance. US Military Academy, West Point, NY. June 17-19, 2002. pp. 17-22.

[2] S. Cheung, K. N. Levitt. “Protecting Routing Infrastructures from Denial of Service Using Co-operative Intrusion Detection,” Proc. of the New Security Paradigms Workshop, Cumbria UK. 1997.

[3] Fred Cohen & Associates. “A Note on Distributed Coordinated Attacks.” April 1996. http://www.all.net/books/dca.

[4] B. Ganter, R. Wille, “Formal Concept Analysis: Mathematical Foundations.” Springer, 1999. ISBN 3-540-62771-5.

[5] X. Geng, A. B. Whinston. “Defeating Distributed Denial of Service Attacks.” IT Pro. July/August 2000. pp. 36-41.

[6] D. W. Gresty, Q. Shi, M. Merabti. “Requirements for a General Framework for Response to Distributed Denial-of-Service.” Seventeenth Annual Computer Security Applications Conference, 10-14 Dec. 2001, New Orleans, LA, USA. pp. 422-9.

[7] J. D. Howard, T.A. Longstaff. “A Common Language for Computer Incidents.” Sandia National Laboratories Report SAND98-8667, October, 1998. Albuquerque, NM.

[8] F. Kargi, J. Maier, M. Weber. “Protecting Web Servers from Distributed Denial of Service Attacks.” WWW10, May 1-5, 2001. Hong Kong. pp. 514-24.

[9] J. Leiwo, T. Aura, P. Nikander. “Towards Network Denial of Service Resistant Protocols.” Information Security for Global Information Infrastructures. IFIP TC11. Sixteenth Annual Working Conference on Information Security. August 22-4, 2000, Beijing, China. pp. 301-310.

[10] J. Mirkovic, Janice Martin, Peter Reiher. “A Taxonomy of DDoS Attacks and DDoS Defense Mechanisms.” Technical Report #020017, Department of Computer Science, UCLA.

[11] S. Mohiuddin, Shlomo Hershkop, Rahul Bhan, Sal Stolfo. “Defending Against a large scale Denial-of-Service Attack.” Proc. of the 2002 IEEE Workshop on Information Assurance. US Military Academy, West Point, NY. June 17-19, 2002. pp. 139-146.

[12] K. Park, H. Lee. “On the Effectiveness of Route-Based Packet Filtering for Distributed DoS Attack and Prevention in Power-Law Internets.” SGICOMM’01, August 27-31, 2001, San Diego, CA. pp. 15-26.

[13] SANS Institute. “Consensus Roadmap for Defeating Distributed Denial of Service Attacks.” Version 1.10. February 23, 2000. http://www.sans.org/ddos_roadmap.htm.

[14] C. Shields. “What do we mean by Network Denial of Service?” Proc. of the 2002 IEEE Workshop on Information Assurance. US Military Academy, West Point, NY. June 17-19, 2002. pp. 196-203.

[15] H. Varian. “Managing Online Security Risks.” Economic Science, Column, The New York Times, June 1, 2000. http://www.nytimes.com/library/financial/columns/060100econ-scene.html.

[16] J. Yan, Stephen Early, Ross Anderson. “The XenoService—A Distributed Defeat for Distributed Denial of Service.” Information Survivability Workshop 2000. Boston, MA.

[17] J. Yan. “Denial of Service: Another Example.” 17th International Conference on Information Security (IFIP/SEC 2002), Cairo, Egypt, May 2002.

[18] V. Gupta, S. V. Krishnamurthy, M. Faloutsos. “Denial of Service Attacks at the MAC Layer in Wireless Ad Hoc Networks”, Milcom 2002, Anaheim, CA (to appear).

[19] Sheriff, L., “Virus Launches DDoS for Mobile Phones,” http://www.theregister.co.uk/content/1/12394.html

[20] L. D. Stein and J. N. Stewart. “The World Wide Web Security FAQ,” Version 3.1.2, February 4, 2002. http://www.w3.org/Security/Faq/.

[21] A. S. Tannenbaum. Computer Networks, Third Edition. Prentice Hall PTR, New Jersey, 1996.

[22] M. D. Torgerson and B. P. Van Leeuwen, “Difficulties with Authenticating Routing Information in Wireless Ad Hoc Networks.” Proceedings of the 2002 IEEE Workshop on Information Assurance, United States Military Academy, West Point, NY. June 2002.

[23] M. Faloutsos, P. Faloutsos, C. Faloutsos. “On Power-Law Relationships of the Internet Topology.” Proc. of ACM SIGCOMM, pp. 251-262, 1999.

APPENDIX A

A Related Work

We studied a representative sample of the recent literature on DoS[5]. Papers from this literature tend to fall into two general categories.

1. Papers that are primarily concerned with ways to define DoS.
2. Papers that are primarily concerned with ways to address the problem.

The definition papers fall into two subcategories.

1a. Taxonomic papers.

1b. Descriptive papers.

The papers focused on addressing the DoS problem fall into four subcategories.

2a. Papers that propose an architecture.

2b. Papers that work on an entire class of attacks.

2c. Papers that provide practical steps.

2d. Papers that present general principles.

A categorization is shown in Figure A-1. The leaves of the categorization tree provide sample papers representing a particular category.

[5] At the coarsest granularity there are two types of DoS attacks—those that do not use the Internet and those that do. In papers written before the Internet (e.g., Gligor, Millen), "DoS" refers to the former set of attacks; in subsequent papers, "DoS" refers to the latter set of attacks, since, by then, the former set of attacks seemed uninteresting. In an effort to distinguish acronymically between these two, some authors have used "DDoS" (D ≡ distributed) (e.g., Geng, Gresty, Mirkovic, Park & Lee, and Yan). Others have used "NDoS" (N ≡ network) (e.g., Shields, and Leiwo) to denote attacks on the network itself. (Kargi uses and distinguishes between DoS and DDoS.) We do not change the acronyms used in the passages we quote from papers. Cohen uses the term "Distributed Coordinated Attacks" (DCAs) for what appears to be a subset of DoS.

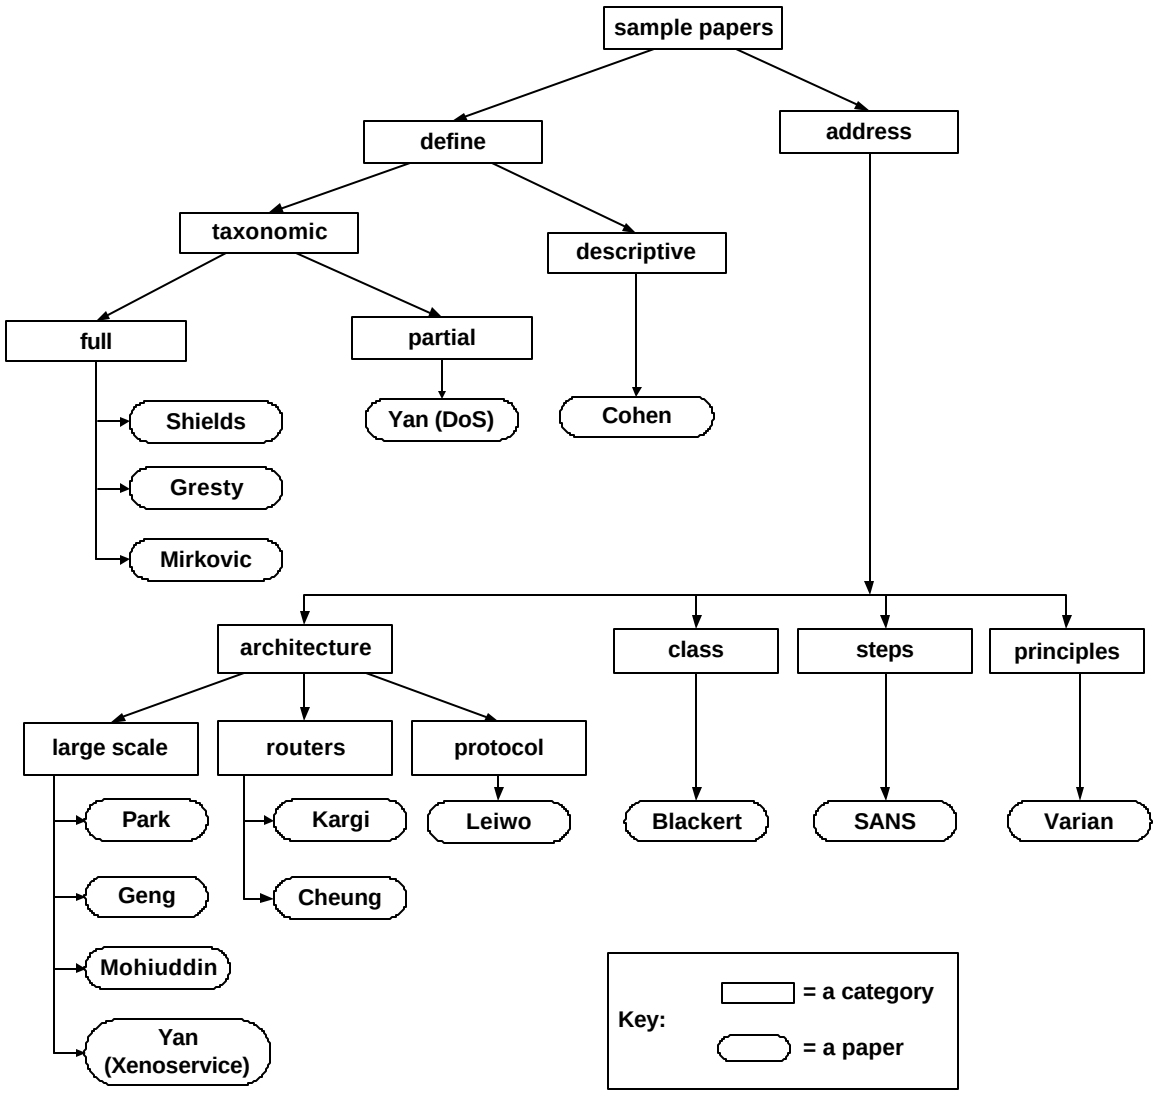

Figure A-1 Sample DoS/DDoS Papers.

The full taxonomic definitions of DoS suggest that it is not yet clear even how to fully define this type of attack. Computer security has been focused on confidentiality and integrity, with little thought, until recently, about availability. Shields [14] builds on DoS research on standalone systems. Gresty et al. [6] provide breadth at the expense of depth, and Mirkovic et al. [10] attempt a comprehensive attack and defense taxonomy.

The large scale architecture papers provide a sampling of the avenues being pursued. Park & Lee [12] capitalize on the “centeredness” of the Internet to limit spoofing. Geng & Whinston [5] leverage economics. Mohiuddin et al. [11] use a separate network to flag DoS traffic, and Yan et al. [16] multiply services.

A promising approach is the one described by Park & Lee [12] on Distributed Packet Filtering (DPF). Building on recent research into the structure of the Internet, DPF is both proactive and reactive. Varian’s newspaper article [15] about insurance complements Park & Lee’s paper, suggesting that we may find solutions in that combination.

We sampled a number of papers on which we present reviews in this section. The papers are presented in reverse chronological order of publication date. We include several comments at the end of the review of each paper, intended only to provide the reader with a counterpoint.

In the References section, we have provided information on several papers that do not appear in the body of this report, namely CERT, Bradley, Dittrich, Ferguson, and the seminal paper by Gligor, Millen, and Yu.

A.1 Blackert et al. [1], "Analysis of Denial of Service Using An Address Resolution Protocol Attack"

The intent of the paper by Blackert et al. is to demonstrate how to address an entire class of attacks, instead of just a single attack. Blackert et al. analyze the ARP Cache Poison attack. They use a modeling and simulation approach, as opposed to either "simplified models" or "testbed analysis." They identify three parameters for this attack: "ARP cache timeout, attacker delay, and network delay." (pg. 19) The first parameter is more important than the other two for effective defense. The simulations show threshold values for the parameters and thus provide a "better understanding" of this particular attack and all of its variants, i.e., the whole class.

The same laboratory, JHU/APL, has provided a similar analysis of other DoS attacks, namely "TCP Syn Flood, Octopus, User Datagram Protocol (UDP) Storm, Snork." (pg. 17). Along the way, Blackert et al. note that there are three general types of DoS attacks (pg. 17):

- software implementation flaws,
- "exploitable protocol designs," and
- "brute force attacks."

Comments:

> Blackert et al.'s analysis gives us the understanding of the parameters at our disposal in addressing this particular class of attack. Presumably, if we can detect an attack of this class, we can use the results of this paper to take effective steps against it, namely reducing the cache timeout until the attack is precluded. At some future time we would restore the timeout to the pre-attack level. However, is it not possible that continuing to lower the cache timeout may lead by itself to a DoS effect?
>
> It appears that Blackert et al.'s approach can be applied to "exploitable protocol designs" only.
>
> Are there non-exploitable protocols? Are there defenses against "software implementation flaws" and "brute force attacks?"

A.2 Shields [14], "What do we mean by Network Denial of Service?"

Shields is clear about his intent:

> The purpose of this paper is to propose a definition of what constitutes a network denial-of- service attack and to put forth for consideration a simple taxonomy of denial-of-service attacks that both includes known attacks and shows where new attacks may be discovered. (pg. 196)

Gligor's definition, which is the starting point of Shields', has three components: Mean Waiting Time, services, and authorized users. Shields' defines "authorized users" as anyone who can send and/or receive network packets, but otherwise his definition is the same as Gligor's:

> A network denial-of-service attack occurs when some set of network entities intentionally uses network services with the goal and effect of causing consumption or corruption of network resources in such a way that some other set of network entities have their ability to access otherwise usable network services degraded or so delayed as to render them unusable (pg. 199).

Shields presents a simple, two part network model consisting of (a) three layers (application, transport, and network), and (b) four node types (host, router, switch, and firewall), where each node has five types of resources available to it: memory (volatile), processing, storage (non-volatile memory), state (protocol code and operating state), and variables[6]. With this definition and model in hand, Shields provides a taxonomy based on a 4-tuple (pg. 200).

1. means ("the general method the attacker uses to cause the attack"),
2. effect ("in general, this will either be to corrupt or consume the resource in question;" corruption involves "crashing" or "conditioning," the latter defined to be the feeding of incorrect learning examples to machines that "learn behaviors or detect anomalies"),
3. resource (see above),
4. location ("the particular network device where the resource being effected exists").

The taxonomy is illustrated by a table with some of the cells filled in with attacks and some cells left empty. The empty cells indicate no known attack. Some of these empty cells could indicate no possible attack, though there is nothing in the taxonomy that would reveal that. The intent of the taxonomy is to facilitate understanding of relationships between attacks.

Comments:

> If Shields' taxonomy were represented by a cube, then each dimension of the cube could represent one of the elements of the 4-tuple. There would be two elements in the "effect" dimension—corruption and consumption—and four in the "resource" dimension—memory, processing, state, and variables. Unfortunately, it is not clear how many elements there would be in the "means" and "location" dimensions. With a specified number of elements for each dimension we could determine the number of cells in the cube and begin to grasp the problem and have a way to search for a solution. Otherwise, we are left dangling.

A.3 Mohiuddin et al. [11], "Defending Against a large scale Denial-of-Service Attack"

Mohiuddin et al. suggest what they believe is a "novel approach for detection and response" to DoS. Their contribution is "an architecture and a method to detect and alleviate denial of service attacks using signature-based methods on a time-partitioned block of data." (pg. 146)

Mohiuddin et al.'s scheme requires its own network, separate from the target operational network. This limits its applicability to subnets. This separate network consists of two types of agents: a set of "sensor agents" (called Data Collection Agents (DCAs)) and a single "managing agent" (called the Data Fusion Agent (DFA)). For each server there is a DCA wrapper that monitors incoming traffic only. The DFA is created by a DCA, whereupon, after a time, it apparently terminates itself, whereupon, sometime later, another DCA creates the DFA. If more than one DFA is in existence, the two creating DFAs decide which one survives. In this way the DFA "moves" around the network and is not an attractive point of attack. Note that the separate network precludes it from DOS attacks due to load on the operating network.

The DCAs send raw data to the DFA. The DFA uses this data to "generate traffic models based on current network traffic" which the DCAs use to detect "anomalous traffic." (pg. 140) Depending upon the "security mode"—passive, suspect, or active—that is set by the DFA, the DCAs drop anomalous traffic or direct it to other servers in order to balance the load.

The DFA works on data in blocks of time. A small percentage of the actual traffic is collected for examination by the DFA. In "suspect" security mode this percentage is higher, and in "active" security mode this percentage is higher still. It appears that one of the benefits of this time block

[6] For processing requests.

approach is that it allows the DFA to operate without being required to preserve state between invocations.

The DFAs use "data mining applied to intrusion detection" to identify the anomalous traffic. This is based on the assumption that there are "common traffic patterns in DOS attacks that differentiate them from normal traffic and other types of attack." (pg. 144) For training data, Mohiuddin et al. used the 1999 DARPA Lincoln Labs data set. Their experiments showed 68% accuracy in identifying "blocks" (i.e., packets).[7] Their experiments confirm their assumption, at least for that particular data set.

Comments:

> While their results are impressive, it is not clear that the effectiveness of a DoS attack could not be maintained simply by generating three times the number of packets. It is unclear what "data mining applied to intrusion detection" is. How would SYN flood be recognized as "anomalous," for example?
>
> Their approach to protecting the centralized service appears to be unique.

A.4 Yan [17], "Denial of Service: Another Example"

Yan discusses DoS attacks that Gresty (Section A.5) classifies as the Consumer type. This type does not appear to be included in any other definition we reviewed, except possibly Kargi's (Section A.7). A simple example of this type of attack is a router that acts as a firewall in that it filters incoming traffic from certain addresses (i.e., it blacklists). This may be welcomed by those on the blocked side of the router, but this could also be considered a DoS attack. Note that there is no exhaustion of resources or bandwidth flooding involved. Resource allocation, stateless protocols, client puzzles are all ineffective against this type of DoS attack.

Yan presents a four part, "simple framework" of DoS "threat models" (pg. 3):

1. Traditional DoS Model: access to one service is blocked to all clients.
2. User-focused DoS Model: access to one service is blocked to a subset of clients.
3. Service-focused DoS Model: access to a set of services is blocked to all clients.
4. Hybrid DoS Model: access to a set of services is blocked to a subset of clients—this is a combination of models 2 and 3.

Yan's blacklist attacks are examples of the Hybrid DoS model. His solution is to "fool or wear off the blacklisting mechanism" by implementing "unblockable Internet services" (pg. 4). He suggests four "possible schemes:"

1. HTTP redirect;
2. Dynamic DNS;
3. Volunteering relay; and
4. Virtual proxies with randomized IP addresses.

The HTTP redirect scheme requires an "intelligent portal server" that can redirect requests to another site so that the returning packets have a different, and, we presume, unblacklisted IP address. The dynamic DNS scheme enables the IP address associated with a DNS name to be changed as needed, depending on load. This is similar to the HTTP redirect scheme except that

---

[7] Accuracy here is defined as (TP+TN) / (TP+TN+FP+FN), where TP and TN are the number of good and bad blocks, respectively, that are correctly identified, and FP and FN are the number of good and bad blocks, respectively, there are incorrectly identified.

the server is distributed. Both of these approaches are resource intensive and, as Yan points out, are subject themselves to the type of DoS attack Yan presents[8]; the blocking filter just has to block a few more IP addresses in order to succeed.

The volunteering relay scheme involves proxy servers that the actual server enlists to respond to queries. Because there can be many proxy servers, each with its own IP address, the service cannot be blocked based on IP address, or at least not until the blacklister has encountered and blocked some critical percentage of the set of proxy servers. This can be a cheap solution.

The virtual proxies with randomized IP addresses scheme is similar to the volunteer relay scheme except that the IP address of the proxy server is randomly chosen by the web server each time the web server receives a request. If the pool of possible proxy server addresses is big enough, the blacklister is defeated.

Comments:

> The importance of this paper is that it fleshes out Gresty's Consumer type DoS attack. But is it really an "attack?" After all, people spend money to be "attacked" in this manner when they pay for the filtering of "adult" content. But one need only look at politically repressive regimes to realize that in that setting this is an attack in its own right.

A.5 Gresty et al.[6], "Requirements for a General Framework for Response to Distributed Denial-of-Service"

Gresty et al. define DoS as follows:

> Denial of service is essentially the problem of an entity within a system (e.g., a user), preventing authorised entities (e.g., other users or programs) having access to resources (e.g., data files, program or network connections) held within the system. (pg. 1)

They state that

> The primary issue for the Internet is trust..." [because]…effective operation of the [Internet] requires that every party behave fairly in every transaction. Any party therefore misbehaving can cause the distributed system to be exploited to deny service to any other party. … Therefore there can be no true technical solution to the Denial of service problem... (pg. 3)

Gresty et al. note that since the Internet cannot guarantee service, DoS in the Internet is an "intractable problem." Preventing DoS is thus an "unreasonable goal" (pg. 5); we should focus instead on response. As a result, they suggest we change the triplet we use to define computer security from "confidentiality, integrity, and availability" to "confidentiality, integrity, and acceptability."

Gresty et al. present a simple classification of denial of service, decomposing it into two problems:

```
Consumer problem (also known as Man-in-the-Middle), and
Producer problem.
```

Implementing a DoS attack via the Consumer problem requires the attacker to get "between" the client and the server. The services provided to the client by the server could then be denied. Gresty et al. note that this approach is not used for well-known DoS attacks. (Yan's "blacklist DoS attacks" (Section A.4) are of this type.)

[8] It is not clear, then, why Yan lists these as "possible" schemes.

Implementing a DoS attack via the Producer problem requires the attacker to produce or offer so many "services" that the recipient is overwhelmed. For example, a SYN is a connection request, which could be seen as an offer of a service, "...and if [the server] has to manage a large number of 'made available' resources this can force it to use up all the internal resources..." (pg. 4) This type of attack is "synonymous with modern DDoS" (pg. 5) since "The nature of the Internet does not allow the victim choice about what it receives..." (pg. 7)

One insight from this paper is that we may be able to find effective responses by using the same principles underlying the nature of the problem:

> If a network incident is not caused by 'normal' random anomalous behavior, but is rather a coordinated hostile incident, it is conjectured that the only effective response is for the entities to deliberately co-ordinate the stabilizing procedure." (pg. 6)

A "general framework for responding" to DoS, given that prevention is intractable, would include the following:

- generality;
- unexploitability; and
- policy (particularly acceptability, but also timeliness, survivability/adaptivity, scalability, and clarity of how to deal with hostile incidents).

Unfortunately (and without surprise), "it can be shown that no existing technique or system currently can meet all these requirements..." (pg. 5) Research issues include (1) deriving "Policy Primitives" to represent state, (2) analyzing network behavior and traffic in detail, and (3) integrating "stabilising techniques" .

Gresty et al. review several papers that we also review. They think that Cohen (Section A.15) identifies the "most important issue" in DoS attacks, namely trust. They express no opinion about the suggestion of Geng & Whinston's (Section A.10) for "economic and technical incentives." They note that the Internet Society's recommendation of ingress & egress filtering to remove from the network packets with spoofed addresses would not stop DoS attacks because some attacks do not require spoofed addresses. The UC Davis' cooperative intrusion detection model is "very important," in the opinion of Gresty et al., because it shows that a coordinated attack requires a co-operative response, which is in tune with their belief that prevention will be built on the same principle as the attacks. However, they think that the UC Davis' WATCHERS protocol, though "perfectly acceptable for a single type of protocol," is not acceptable for "general transactions" and thus is suitable only for a "restricted" environment. And finally, they note that Gligor's paper shows that using "maximum waiting time" makes DoS a simple matter to detect but that "without an alternative route to request resources or direct requests then there is no way to do anything other than detect the denial of service." (pg. 3)

Comments:

> Dividing DoS attacks into Producer and Consumer attacks seems to provide little insight. However, this is the only definition that includes Yan's "blacklist DoS attacks" (Section A.4) so perhaps it provides more insight than is apparent.
>
> If we could develop a general framework that was "unexploitable," then we should be able to bootstrap to an entire system that is unexploitable, thereby preventing DoS and contradicting one of their conclusions.
>
> If the fundamental problem or issue with the Internet is "trust," then the fundamental question, we think, is, Can that infrastructure be made trustworthy? or, similarly, Do all of the people using the infrastructure have to be trustworthy before we can expect relief from DoS attacks?

A.6 Park & Lee [12], "On the Effectiveness of Route-Based Packet Filtering for Distributed DoS Attack and Prevention in Power-Law Internets"

Park & Lee address IP spoofing. This, in their opinion, is the key to DoS prevention:

> As with prank telephone calls or ringing of door bells in days gone by, an effective means of preventing DoS attacks from occurring in the first place—also the only fundamental solution given the intrinsic susceptibility of service provisioning systems to DoS—lies in identification of the attacker which admits assigning commensurate costs (e.g., legal or economical) to the perpetrating entity. (pg. 15)

Park & Lee's "main contribution" is "advancing a scalable architecture for DDoS attack prevention that is effective for Internet AS [autonomous system] [9] topology." (pg. 16) Their architecture uses route-based filtering in a scheme they call "distributed packet filtering" (DPF). DPF is designed to run on Internet routers. It depends "intimately on the connectivity structure of the underlying AS graph" of the Internet. In particular, we show that power-law Internet AS topology is crucial in facilitating small coverage with strong proactive and reactive filtering effect." (pg. 16) "Power-law graph structure induces 'centers' where connectivity is concentrated on a few nodes..." (pg. 20) Park & Lee explain the approach with the following analogy:

> This is akin to setting up road blocks [DPF filters] at certain intersection points in a city to apprehend bank robbers [DoS traffic]: the bank robbers are constrained to take the public transportation system whose routes, in turn, are constrained by the physical street network (topology) and routing policy/algorithm imposed by the municipal transportation department. (pg. 16)

And this is Park & Lee's description of DPF's "main strength:"

> Route-based DPF's main strength lies in the fact that with partial coverage or deployment— about 18% in Internet AS topologies—a synergistic filtering effect is achieved whose collective filtering action proactively prevents spoofed IP flows from reaching other autonomous systems in the first place." (pg. 16)

Park & Lee argue that "eliminating all spoofable IP flows is an unrealistic goal given its inherent difficulty with respect to Internet AS [autonomous system] connectivity." (pg. 22) However, with only about 18% coverage, the number of IP spoofed packets that reach their destination is "sparse," and, "the attacker's source AS's can be localized to within 5 sites (a constant) independent of system size." (pg. 16)

The goals for DPF are fourfold: (1) maximize the percentage of spoofed packets that are filtered, (2) minimize the number of possible attack-generating sites for the packets that are not filtered, (3) minimize the number of sites that filter, and (4) find the optimal location for filtering sites. (pg. 18) A "maximal" route-based filter requires $O(n^2)$ space—an "overwhelming burden" on routers. A "semi-maximal" filter requires only linear space but it is only "marginally" less capable than a maximal filter, making it feasible, if not practical. Placing semi-maximal filters at those "'centers' where connectivity is concentrated" accomplishes the four goals of DPF.

Park & Lee review a number of related techniques, including (a) IP traceback techniques such as "link testing" (slow, manual, with high overhead), "audit trail" (high overhead at routers), and "behavioral monitoring" (the criminal must return to the "scene" in order for this to work), (b) "packet-based traceback," and "information packets." Park & Lee's work appears in closest

---

[9] According to Faloutsos [23], "The definition of an autonomous system can vary in the literature, but it usually coincides with that of the domain." Faloutsos defines a "domain" to be a subnetwork under "separate administrative authorities."

competition with "probabilistic packet marking" (PPM). The drawbacks of PPM are that it is reactive only, it requires header bits, and its uncertainty is proportional to the number of attack hosts. Park & Lee show that DPF is superior to PPM on all three counts.

Comments:

> With additional research (and corresponding results) and used in conjunction with Varian's type of suggestion (Section A.11), this approach suggests solutions.
>
> In the meantime, how much of the spoofed flow does 18% coverage filter? Apparently it is not a constant, independent of network size, which would be ideal. Nor is it a percentage of the flow, which is second best. How many is "sparse?"
>
> The attackers source AS's can be localized "to within 5 sites..." What is a "site?" Is it one machine at one IP address? Is it a class C or B or A address? Is it a router or a "domain" (which we take to mean a set of routers)?

A.7 Kargi et al. [8], "Protecting Web Servers from Distributed Denial of Service Attacks"

Kargi et al. present a method of protecting web servers from DoS attacks, using upstream control based on traffic monitoring and class based routing. They have tested their approach with a testbed that has the following components: (a) eight "attack" clients, (b) one normal client, (c) a load balancer, and (d) three web servers. Incoming packets (i.e., incoming to the web servers from the attack and normal clients) go through the load balancer; outgoing messages bypass the web servers. The software "Traffic Monitor" on the load balancer partitions the incoming traffic into four "classes," for each of which there is a separate queue, using the Linux "Class Based Queueing" (CBQ). The attack machines can generate "about 8 times the input capacity of the load balancer," but the balancer is able to continue almost-normal service for the normal client by relegating incoming attack streams to classes that provide lower levels of service. The attacks used in this testbed include SYN floods, HTTP requests, Smurf attacks, TARGA flooding (using TFN2K, which is related to TFN). Kargi et al. show the limitation of this approach, namely that the Traffic Monitor must be able to sense an attack in order to be able to defend against it.

Kargi et al. provide a simple classification of attacks, based on three categories:

1. *System attacked*: firewall, router, load-balancer, web server, or "supporting services" (such as a database server);
2. *Part of the system attacked*: hardware (a "rare" occurrence), link, operating system, network stack, application;
3. *Bug or Overload.*

They place "common" DoS attacks into this classification to show that the classification is "suitable for distinguishing" attacks (pg. 516).

Kargi et al. argue that there is "no complete solution" yet for DoS attacks. Nevertheless, there are "basic security measures," which they consider "mandatory" and would, presumably, preclude a number of attacks. They note that "many experts think that the only durable solution" is to "improve the security on all hosts." (pg. 519) While this may be the case, Kargi et al. do not think that improving security across the board will happen in the "foreseeable future." And even if it does, the continuing growth of the Internet "will simply absorb" the benefits. Nor do they think that usage-based approaches, such as Geng & Whinston (Section A.10) describe, will be adopted.

Kargi et al. think that the next step in attacks may be "to automate the process of intruding daemons which again spread themselves to other hosts." (pg. 518)

Comments:

> This is a purely reactive approach. It is not clear how protected the load balancer is from becoming the bottleneck. If more clients were to join in the attack and generate not just 8 times the input capacity of the balancer but, say, 80 times the capacity, would the load balancer be able to keep up?

A.8 Yan et al. [16], "The XenoService – A Distributed Defeat for Distributed Denial of Service"

Yan et al. describe a "workable solution" in what they call a "distributed defeat for distributed denial of service." That is, one way to defeat distributed DoS attacks is by distributing the service that is being attacked. If sharing a resource enables the problem, then sharing a resource should be able to enable the solution. This is an example of an insight noted in Gresty et al.'s paper (Section A.5).

A XenoService is a network of web hosts that "respond to an attack on any one web site [in the group] by replicating it rapidly and widely." As a result, the resources that an attacker must expend to defeat any member of a XenoService is dramatically higher than it would be if the victim were not a member of the XenoService. XenoServers run on Nemesis, an operating system that provides quality-of-service guarantees and thus may be able to defeat resource exhaustion type of DoS attacks.

Yan et al. assert that the "widely believed" ways to prevent DoS attacks, namely secure hosts, egress filtering, and fixes for specific vulnerabilities, are not enough. DoS is a "system problem," not due to a specific technology: secure systems (i.e., ones that are "Orange-book- evaluated") are expensive and unavailable, and there is no economic incentive to prevent DoS attacks. DoS is an example of the "Tragedy of the Commons:" "…while everyone may have an interest in protecting a shared resource (Internet security), individuals have a strong motive to cheat (connecting insecure computers)." (from the Abstract) Any solution must include economic forces, which concurs with Varian's view.

Comments:

> This approach makes sense. But is it possible that a DoS attack could use XenoServers to amplify the attack, thereby engulfing the entire set of Xenoservers?

A.9 Leiwo et al. [9], "Toward Network Denial of Service Resistant Protocols"

After reviewing general work on DoS, attacks, and protocol design principles, Leiwo et al. conclude that "only limited protection of availability can be achieved through technical measures," (pg. 301) and that "any statefull handshake protocol" is vulnerable to SYN flooding types of attacks (from the Abstract). They note that the 1985 DoS workshop concluded that "no generic, mission independent DoS conditions can be identified" (pg. 302) and that The Committee on Information Systems Trustworthiness, as recently as 1999, determined that "neither DoS prevention methods nor systematic design methods exist against DoS." (pg. 302) However, the authors state that there are a number of ways of "dealing" with network DoS (pg. 302), which we presume means that while we cannot prevent DoS attacks we can take some steps against them.

Leiwo et al. arrange "attacks and attack methods" into the following classification:

```
Tolerable attacks (preventable by proper protocol design)
Normal protocol execution only
Deviation attacks
```

```
Deviation from protocol sequence (e.g., SYN flood)
Deviation from protocol syntax (e.g., Ping attack)
Deviation from protocol semantics (e.g., IP spoofing)
Fabrication attacks (e.g., forging routing information)
Fatal attacks (e.g., server intrusion)
```

Leiwo et al. then present several DoS-resistant protocol design principles:

- allocating memory only after authentication,
- exhaust the cheapest detection techniques first,
- parameterize the work load of the client and enable the server to modify the client work load so the server can react to DoS (that is, if the server suspects client A of mounting an attack, the server can increase the workload required of client A to engage in the protocol).

Leiwo et al. present an example "network DoS resistant variant of the X.509 authentication protocol" using their DoS-resistant design principles.

Comments:

> While it is clear that there are steps we can take to make DoS more difficult to achieve, it is not clear how significant those steps are. For example, if we were to implement the authors' example protocol, would it require only twice as many zombie machines to produce the same DoS effect? Even if it were to require at least six orders of magnitude more zombies, would this do more than buy a little time? But then, perhaps buying time is all we can expect, or all that we need?

A.10 Geng & Whinston [4], "Defeating Distributed Denial of Service Attacks"

Geng & Whinston argue that "Globally coordinated solutions are indispensable for defeating the DDoS attack. Fostering such solutions will require proper economic incentives for all parties affected (directly or indirectly) by DDoS attacks." (pg. 41) Geng & Whinston propose what they call an "e-postal service framework."

DoS exploits (1) insecure channels and platforms, (2) easy traffic volume generation, (3) victim circumvention (e.g., clog paths to a victim, rather than the victim himself), and (4) identity hiding. Geng & Whinston think that each of these must be addressed in order to counter DoS. The possible solutions are presented in a simple, two-tiered taxonomy consisting of "local" and "global" solutions." For each local solution, Geng & Whinston point out the limit of its effectiveness (included below in parentheses). With all of the local solutions discounted, only global solutions remain as possible.

```
Local solutions
    Local filtering
        (Effective only if the attacker cannot clog the filter
          and thus continue the attack that way.)
    Changing IPs
        (Effective only until attackers do not attack based on IP
          addresses.)
    Creating client bottlenecks (e.g., client puzzles)
        (Effective only if the client needs to solve the puzzle
          in order to continue the attack.)
Global solutions
    Improving the security of the entire Internet
    Using globally coordinated filters
    Tracing the source IP address
```

The global solutions only work if everyone does them, so it is not in anyone's interest to implement the solutions unless everyone else does too. The problem, as Geng & Whinston see it, is that users do not pay based on the amount of traffic they generate. Rather, users pay a "flat" (i.e., usage-independent) fee. If usage-based fees were in place, then individual users would have an economic incentive to prevent their machine from participating in a DoS attack. (An alternative solution would be to arrange to have one organization own enough of the Internet that any DoS attack would impact its profits. Such an organization would have the power to act and the incentive to do so. Unfortunately this would also encourage a monopoly, so Geng & Whinston dismiss it.)

Geng & Whinston recommend the adoption of an "e-postal service framework." The framework consists of a proper subset of the Internet routing infrastructure, referred to as a "e-postal system." It is not clear how membership in this subset would be determined; presumably each router owner would be free to make the decision and would agree to the rules specified by whatever organization it is that runs the framework. At any rate, routers are either entirely in the e-postal system or entirely outside the e-postal system: they cannot route both secure and insecure traffic.

The routers on the frontier of the framework, the "edge routers," have two functions: "store and issue e-stamps" and "authenticate users." The stamps cost money, thereby creating the incentive for those using the e-postal system to avoid participation in DoS attacks. It is not clear how the authentication works or what it is intended to do. Perhaps each stamp has a name associated with it and this serves for the authentication. Presumably the e-postal system is so large and with so many edge routers that it would not be possible to mount a successful DoS attack against it, such as flooding the edge routers with requests for e-stamps.

This is a "usage-based" scheme: someone pays for all traffic. The assumption is that DoS attacks require a small number of servers to generate a high volume of packets. If the owners of the servers have to pay for usage, then they will by themselves shut down such high volume.

Comments:

> The e-postal service framework suggests a single, controlling organization, beyond the oversight of any national government, with lots of money changing hands and lots of opportunities for the abuse of power, the manipulation of the flow of information, and the invasion of privacy.

A.11 Varian [15], "Managing Online Security Risks"

Hal Varian, in an article in *The New York Times* , presents the economic argument for better risk management in the Internet. He states that "...one of the fundamental principles of the economic analysis of liability [is that] it should be assigned to the party that can do the best job of managing risk." Once that assignment is made, the party will want to buy insurance. The insurance companies will then be in a position to benefit economically from providing good guidelines on security; the better the guidelines, the less they have to pay in claims, the lower their rates, and the higher their market share. This will drive the insurance companies to oblige those in control of the Internet to meet security risks.

Comments:

> We believe that this model does work but only when, for each person, increasing that person's security decreases that person's residual risk. But this characteristic does not hold for all DoS attacks. However, if we were able to proactively blunt DoS attacks and, simultaneously, trace the attacks to a small set of machines, as Park & Lee's research (Section A.6) is pursuing, then we

believe that Varian is on the right track.

A.12 SANS Institute [13], "Consensus Roadmap for Defeating Distributed Denial of Service Attacks"

SANS Institute published a report in 2000 written by 17 experts on how to defeat DoS attacks. The document includes the following sections:

- "Key Trends and Factors" about current DoS attacks,
- "Immediate Steps to Reduce Risk and Dampen the Effect of Attacks," and
- "Longer Term Efforts to Provide Adequate Safeguards."

The Key Trend, if there is a single one, is that attacks will be getting more virulent, for which "broad community action" is required.

The "Immediate Steps" section outlines four problems:

1. "Spoofing;"
2. "Broadcast Amplification" (also known as Smurf attacks);
3. "Lack of Appropriate Response to Attacks;" and
4. "Unprotected Computers."

For each of the problems, "solutions" are listed. It is hoped that "major users" will set the example by adopting the solutions listed.

The "Longer Term Efforts" section includes research in "areas of policy and law to enable us to deal with aspects of the problem that technology improvements will not be able to address by themselves." (pg. 6) That is, technology alone cannot solve this problem. As the authors put it, their hope is to "...manage the emerging risks and keep them within more tolerable bounds."

This is intended to be a "living" document, meaning that SANS plans on updating it in the future. It is also a "consensus" document, meaning that it is a "product of the joint thinking of some of the best minds in security." (pg. 7)

Comments:

> This is a reasonable short-term approach to the problem, to keep it within "tolerable bounds." Perhaps instead of using the word "Defeating" in the title, they should have used the word "Addressing." SANS has also produced a "step-by-step" guide with smaller scope.

A.13 Mirkovic et al. [10], "A Taxonomy of DDoS Attacks and DDoS Defense Mechanisms"

The goal of this paper is the imposition of structure on DoS activity, both attacks and defenses, to "foster easier cooperation among researchers." (pg. 11) The two taxonomies are intended to be "useful to the extent that they clarify our thinking and guide us to more effective solutions." Mirkovic et al. are "not aware of any other attempt to introduce formal classification into the DDoS attack mechanisms." (pg. 10)

Mirkovic et al. identify a number of areas of study: protocol security mechanisms, attack isolation strategies, rate limiting mechanisms, overlay networks, pushback, elimination of spoofing, resource accounting, anomaly detection, and traceback.

The attack and defense taxonomies that Mirkovic et al. provide have four levels each, with 4, 9, 5, and 14 items on each level of the attack taxonomy, and 2, 5, 5, and 15 items on each level of the defense taxonomy.

The top level of their attack taxonomy has four items:

- degree of automation (manual, semi-automatic, and automatic),
- exploited vulnerability (protocol, brute-force),
- attack rate dynamics (continuous, variable),
- impact (disruptive, degrading).

The top level of their defense taxonomy has two items:

- activity level (preventive, reactive),
- deployment location (victim network, intermediate network, source network).

Comments:

> DoS attacks exploit cooperation and it is logical to assume that a DoS defense should explore the same principle. Mirkovic et al. have made a step in that direction.
>
> Unfortunately they do not provide us with a clear way to proceed since we cannot use their taxonomies to place attacks or defenses. The taxonomies appear to be closer in structure to a series of categories than taxonomies, in the same spirit as Howard & Longstaff's work.
>
> We do not agree that DoS requires the use of "subverted" machines since reflection attacks are a counterexample.
>
> Also, we do not agree with the statement: "Deploying comprehensive protocol and system security mechanisms can make the victim completely resilient to protocol attacks." (pg. 7) If the design of the Internet makes DoS possible, then do there exist alternative designs that would preclude DoS?

A.14 Cheung et al. [2], "Protecting Routing Infrastructures from Denial of Service Using Cooperative Intrusion Detection"

Cheung et al. present "a first step" in a solution to the problem of routers that misroute packets, including dropping packets that should have been passed on. Cheung et al. are not addressing detection of a SYN flood passing through a router, for example. Rather they are addressing attacks on the network itself—NDoS in the true sense of the acronym.

Participating routers evaluate neighboring routers via test packets. That is, they "cooperatively diagnose each other." Cheung et al. present two protocols: "autonomous distributed probing," and "source-initiated distributed probing." Both protocols are shown to be "sound, complete, and responsive." That is, the protocols have the following properties: "(1) A well-behaved router never incorrectly claims another router as a misbehaving router; (2) If a network has misbehaving routers, one or more of them can be located; (3) Misbehaving routers will eventually be removed." (from the Abstract)

Comments:

> This is a relatively new area of research, even though Cheung et al. speak of the "1980 ARAPNET collapse" that was due to a faulty router. Commandeering a router is an attack on the network itself and potentially powerful. This seems like fertile ground for an asymmetric information warfare attack. However, for the time being, DoS attacks seem to be against sites, not the network. (See the review in the Bibliography of Bradley's paper.)

A.15 Cohen & Associates [3], "A Note on Distributed Coordinated Attacks"

Cohen describes an attack on the allnet site, as well as general information about the set of

DCAs, which appears to denote a subset of DoS attacks, as we have been using the term.[10] DCAs succeed by exploiting "workload advantage," the same principle that cryptography uses. Cohen considers DCAs a "nearly ideal IW [Information Warfare] weapon."[11] Some of the characteristics he notes are indirectness, distributed contribution, unawareness, and "open loop" (apparently meaning that there is no required protocol response in order for the attack to succeed).

Cohen considers that prevention is "unlikely to be fully effective in the near future," since today's environment has "deeply embedded design flaws" (pg. 12). Detection is easy via threshold schemes, such as the detection of three failed login attempts from a given site. "As far as we can tell, the only way to reliably track down a DCA is by coordinating audit trails between a set of sites that form a complete path between the attacker and the target" (pg. 14). (He mentions the use of search engines to track down sources of attacks.) "Based on our results to date, we believe strongly that the most effective DCA defense today is an automated zero- tolerance approach to reporting detected anomalies, and that such a defense will require community tolerance and vigilance." (pg. 19) Cohen also provides a "mathematical characterization" of DCAs.

Comments:

> What are the "deeply embedded design flaws" of the Internet? And, more importantly, are they inherent to the Internet's functionality? In particular, is it possible to design a system with similar functionality that does not have design flaws that would lead to DoS attacks?

---

[10] The "distributed coordinate attack" is Attack 73 in Cohen's list of 94 "information system attacks." Also in the list is "Attack 72: network service and protocol attacks," and "Attack 74: man-in-the-middle." Attack 74 corresponds to Gresty's Consumer type (and Yan's blacklist type of attack). "DoS," as we have been using it in this paper, appears to includes Cohen's Attacks 72-74, making DCAs a proper subset of DoS.

[11] But then, a recent book, "Information Warfare: How to Survive Cyber Attacks" by Michael Erbschloe (McGraw-Hill, 2001, ISBN 0- 07-213260-4), mentions DoS attacks only tangentially.

APPENDIX B

B DoS Attacks mapped to Communication Stack View

The Communication Stack View, presented in Section 3, tracks DoS attacks by their activity, both normal and abnormal, in the communication stack of the Origination, Facilitation, and Victim Networks. Table B-1 presents examples of this approach in a compact manner for a list of attacks. This tabular presentation provides an effective way of displaying a summary of the use of the communication stack by DoS attacks. Appendix C provides a description of the list of attacks found in Table B-1.

Activity for some attacks is grayed-out in Table B-1 because these entries are for tools that have multiple attack methods. For example, many of these tools can launch TCP floods, ICMP floods, UDP floods, Smurf attacks, and/or randomized packets. Therefore, the far left entry for that attack directs the reader to "See <other attacks>" because the traits for all of the available attacks in that tool can't be condensed into a single line, and each possible attack method is already covered in its own line, like "TCP Flood". The reader is directed to see the details for the sub-attacks listed under that attack name.

● = Layer utilized normally, ● = Layer utilized abnormally															
	Origination Network					**Facilitation Network**					**Victim Network**				
DoS Attack	Physical Layer	Data Link layer	Network Layer	Transport layer	Application Layer	Physical Layer	Data Link layer	Network Layer	Transport layer	Application Layer	Physical Layer	Data Link layer	Network Layer	Transport layer	Application Layer
TCP SYN Flood	●	●	●	●		●	●	●			●	●	●	●	
Smurf (ICMP Based)	●	●	●			●	●	●			●	●	●		
UDP Flood	●	●	●[1]	●		●	●	●[1]			●	●	●	●	
ICMP Flood	●	●	●[1]			●	●	●[1]			●	●	●		
ECHO-CHARGEN	●	●	●	●		●	●	●			●	●	●[2]	●	
Ping of Death	●	●	●			●	●	●[3]			●	●	●		
Teardrop	●	●	●			●	●	●[3]			●	●	●		
Land Attack	●	●	●	●		●	●	●			●	●	●[4]	●[4]	
Trin00 or trinoo Command Channel (attacker < --> slaves)	●	●	●	●		●	●	●			●	●	●	●	
Trin00 or trinoo Attack Channel (slaves <--> target) See UDP Flood															
TFN Command Channel (attacker <--> slaves)	●	●	●[5]			●	●	●[5]			●	●	●[5]		

1. Spoofed source IP addresses are sometimes used (which could be detected), but are not required.
2. Some variants spoof SRC IP as the DST IP (detectible at target).
3. If fragment tracking and reassembly is done.
4. It helps if the Network and Transport Layers are looked at in conjunction.
5. Only if "pseudo-stateful" tracking is done (ICMP Pong without an initiating Ping).

Table B-1: Table of DoS attacks mapped to the Communication Stack View.

● = Layer utilized normally, ● = Layer utilized abnormally															
	Origination Network					**Facilitation Network**					**Victim Network**				
DoS Attack	Physical Layer	Data Link layer	Network Layer	Transport layer	Application Layer	Physical Layer	Data Link layer	Network Layer	Transport layer	Application Layer	Physical Layer	Data Link layer	Network Layer	Transport layer	Application Layer
TFN Attack Channel (slaves < -- > target) See TCP SYN Flood, UDP Flood, ICMP Flood, and Smurf															
Stacheldraht[1]															
TFN2000 or TFN2K[2] Command Channel (attacker <--> slaves) See TFN Command Channel.															
TFN2000 or TFN2K Command Channel (attacker <--> slaves) See TCP SYN Flood, UDP Flood, ICMP Flood, Smurf, and TARGA															
TARGA	●	●	●[3]	●[3]		●	●	●[3]			●	●	●[3]	●[3]	
DNS Cache Poisoning	●	●	●	●	●	●	●	●			●	●	●	●	●
ARP Cache Poisoning	●	●									●	●[4]			
802.11 Wireless RTS/CTS Interference	●	●				●					●	●			
Application Buffer Overflows or Unexpected Input	●	●	●	●	●	●	●	●			●	●	●	●	●[5]

1. See Trin00 and TFN for most probable configurations.
2. TFN2K includes various features that are supposed to "hide" the Command Channel.
3. "Randomized" fields in some combination. May or may not be detectible.
4. Rate of ARP-replies might be a clue.
5. If application has any such checks/detection capability programmed in.

Table B-1 (cont.): Table of DoS attacks mapped to the Communication Stack View.

APPENDIX C

C DoS Attack List

In this appendix, we provide a list of DoS attacks. The listed attacks were effective enough to be noticed and studied to some degree by the security community. There are certainly other DoS attacks existing that are not listed, but many of these either have no available data to categorize them, are an early prototype version of one of the listed attacks, are a slight variation on one of the listed attacks, or didn't work in the first place and so were not often employed. Such attacks were not used enough to be noticed/documented or were identified as the most similar attack or a tweak thereon.

Some other DoS incidents not listed are the unintended or incidental results of other actions.

- A firewall or DNS misconfiguration accidentally blackholes your network from the world or vice-versa.
- Incidental resource exhaustion from worm propagation (Code Red, ILOVEYOU, etc).
- The "Slashdot effect" where a site is crashed by the sheer volume of legitimate traffic trying to access a resource.

Approximate Dates are the earliest recorded occurrence found in CERT advisories, SANS articles, or dates of release/detection on white and black-hat sites. They may not be completely accurate. Brief Descriptions are general overviews or summaries of the attack. Where possible, or applicable, a reference link or citation to a more descriptive document is included.

C.1 TCP SYN Flood (half-open connections)

Approximate Date: 1996 (September)

Impact:
- Resource exhaustion
  - Network stack buffers
- Resource “gone”
  - OS crash (rare but possible)

Brief Description: http://www.cert.org/advisories/CA-1996-21.html

The TCP three-way-handshake is initiated, but the final ACK is never sent (usually due to spoofed IP source addresses that are unreachable). By establishing enough half-open connections, the pending connection buffers on the victim are exhausted, preventing legitimate hosts from establishing connections to the victim.

C.2 UDP Flood

Approximate Date: unknown (1999 at the latest)

Impact:
- Resource exhaustion
  - Network stack buffers
  - Application I/O buffers
  - Bandwidth fully utilized

Brief Description: http://rr.sans.org/threats/understanding_DDoS.php

UDP packets are sent to random ports on the victim machine. If the port is open, the crafted packet must be sent on to the application listening on that port. If the port is closed (the more common situation), the victim must generate an ICMP “destination unreachable” message, which consumes further bandwidth and CPU time.

C.3 ICMP Flood

Approximate Date: unknown (1999 at the latest)

Impact:
- Resource exhaustion
  - Network stack buffers
  - Bandwidth fully utilized

Brief Description: http://rr.sans.org/threats/understanding_DDoS.php

A large enough volume of ICMP traffic is sent to overwhelm the victim's resources or consume the available bandwidth. If ICMP "echo request" messages are being sent, more bandwidth is wasted when the victim generates ICMP "echo reply" messages.

C.4 ECHO-CHARGEN

Approximate Date: 1996 (February)

Impact:
- Resource exhaustion
  - Network stack buffers
  - CPU busy
  - Bandwidth exhaustion

Brief Description: http://www.cert.org/advisories/CA-1996-01.html

A self-maintaining flood of traffic is initiated between one or two systems using a combination of the echo and/or chargen ports. A spoofed UDP packet starts the exchange, and then the echo and/or chargen services maintain or escalate the traffic level.

C.5 Ping of Death

Approximate Date: 1996 (December)

Impact:
- Service "gone"
  - OS crash

Brief Description: http://www.cert.org/advisories/CA-1996-26.html

A ping packet that is larger than the maximum size is fragmented and sent to the target. When the target machine reassembles the fragments, an IP datagram larger than 65535 octets is produced. On vulnerable systems, this causes a buffer overflow and results in an OS crash.

C.6 Teardrop

Approximate Date: 1997 (November/December)

Impact:
- Service "gone"
  - OS crash

Brief Description: http://www.cert.org/advisories/CA-1997-28.html

Overlapping IP fragments are sent to the victim. A vulnerable implementation of IP fragment re-assembly will result in an OS crash.

C.7 Land Attack

Approximate Date: 1997 (November/December)

Impact:
- Service "gone"
  - OS crash or "freeze"

Brief Description: http://www.cert.org/advisories/CA-1997-28.html

A crafted packet is sent to an open TCP or UDP port on the victim. The packet's source IP and source port are set to be the same as the destination IP and destination port (i.e., 10.0.0.1:139 to 10.0.0.1:139). A vulnerable machine will crash, "freeze," or lose network functionality.

C.8 Smurf

Approximate Date: 1998 (January)

Impact:

- Resource exhaustion
  - Network stack buffers
  - CPU busy
  - Bandwidth exhaustion

Brief Description: http://www.cert.org/advisories/CA-1998-01.html

Forged ICMP echo requests are sent to broadcast addresses. If the routers are configured to pass broadcast IP addresses, all hosts on the network will reply. This results in an ICMP echo reply flood to the spoofed source address in the initial echo request.

C.9 Trin00 or trinoo

Approximate Date: 1999 (July)

Impact:

- Resource exhaustion
  - Network stack buffers
  - CPU busy
  - Bandwidth fully utilized

Brief Description: http://www.cert.org/incident_notes/IN-99-07.html

Distributed Denial of Service network, which can coordinate UDP floods from multiple compromised systems.

C.10 TFN (Tribe Flood Network)

Approximate Date: 1999 (July)

Impact:

- Resource exhaustion
  - Network stack buffers
  - CPU busy
  - Bandwidth fully utilized

Brief Description: http://www.cert.org/incident_notes/IN-99-07.html

Distributed Denial of Service network, which can coordinate UDP floods, TCP SYN floods, ICMP echo floods, or Smurf attacks.

C.11 Stacheldraht

Approximate Date: 1999 (August)

Impact:

- Resource exhaustion
  - Network stack buffers
  - CPU busy
  - Bandwidth fully utilized

Brief Description:

Stacheldraht is functionally a combination of Trin00 and TFN with encrypted command channels.

C.12 TFN2000 or TFN2K (Tribe Flood Network 2000)

Approximate Date: 1999 (December)

Impact:
- Resource exhaustion
  - Network stack buffers
  - CPU busy
  - Bandwidth fully utilized

Brief Description: http://www.cert.org/advisories/CA-1999-17.html

TFN2000 is an updated and "improved" version of TFN that adds encrypted command channels and other command channel hiding features.

Distributed Denial of Service network, which can coordinate UDP floods, TCP SYN floods, ICMP echo floods, or Smurf attacks.

C.13 TARGA

Approximate Date: unknown

Impact:
- Service "gone"
  - OS crash or "freeze"

Brief Description:

TARGA randomizes some combination of fields in the IP packets it creates. These fields are in the Network and Transport layers and are meant to probe for vulnerable combinations in the target's IP stack. A vulnerable machine will most likely either crash or freeze.

C.14 DNS Cache Poisoning

Approximate Date: unknown (1997 at the latest)

Impact:
- Service "gone"
  - Service hidden/misdirected
  - DNS entry removed/changed

Brief Description: http://rr.sans.org/firewall/DNS_spoof.php

False authoritative DNS records are sent to name servers, which add them to their caches. Queries are then misdirected to non-existent or incorrect IP addresses.

C.15 ARP Cache Poisoning

Approximate Date: unknown (2001 at the latest)

Impact:
- Service "gone"
  - Service hidden/misdirected

Brief Description: http://rr.sans.org/threats/address.php

False IP to MAC address mappings are sent on to a LAN segment. Machines on that LAN will add the mappings to their ARP cache, resulting in IP traffic not being delivered to the appropriate destination. Non-existent or incorrect MAC addresses can be used to deny access to/by a specific machine or for Man-in-the-Middle attacks.

C.16 802.11 Wireless RTS/CTS Interference

Approximate Date: unknown (mostly theoretical/academic exercises at this point)

Impact:
- Resource exhaustion
  - Bandwidth exhausted

Brief Description:

Properly timed RTS (Request to Send) or CTS (Clear to Send) frames are generated with the largest possible NAV (Network Allocation Vector) value in order to reserve the wireless channel. This channel time then goes unused, but prevents other nodes from using the channel. The process is repeated to degrade/deny all legitimate communication on the wireless channel.

C.17 Application Buffer Overflows or Unexpected Input

Approximate Date: continual

Impact: The most common occurrences include any of

- Resource exhaustion
  - Application I/O buffers
  - CPU busy
  - Disk space
- Service “gone”
  - Application crash or hang

Brief Description:

This category includes too many examples to enumerate. In general a buffer overflow or an input parsing error results in an application crash, stalled or runaway processes, recursive creation or traversal of files and directories, and other actions. This behavior has been observed in web servers, WINS name servers, DNS name servers, FTP servers, SNMP servers, game servers (doom and quake), and various network-capable media players. These are generally short lived as they are superseded by access-based exploits using the same flaws (such as gaining root/admin or user access to the machine) and patches follow shortly after that.

DISTRIBUTION:

1 MS 0785
R. E. Trellue, 6501

2 MS 0451
K. L. Shanklin, 6504

1 MS 0784
R. A. Duggan, 6512

1 MS 0784
M. J. Skroch, 6512

1 MS 0784
B. J. Surbey, 6512

1 MS 0784
M. L. Young, 6512

1 MS 0785
T. J. Draelos, 6514

1 MS 0785
T. S. McDonald, 6514

1 MS 0785
T. A. Obenauf, 6514

1 MS 0785
M. D. Torgerson, 6514

1 MS 0785
M. J. Berg, 6516

1 MS 0785
P. L. Campbell, 6516

1 MS 0785
J. D. Dillinger, 6516

5 MS 0785
D. P. Duggan, 6516

1 MS 0785
B. L. Hutchinson, 6516

1 MS 0785
D. Kilman, 6516

1 MS 0785
B. P. Van Leeuwen, 6516

1 MS 0785
W. F. Young, 6516

1 MS 9018
Central Technical Files, 8945-1

1 MS 0899
Technical Library, 9616

1 MS 0612
Review & Approval Desk, 9612
For DOE/OSTI