Detection and Reaction to Denial of Service Attacks

G. Koutepas, B. Maglaris
Network Management & Optimal Design Laboratory
Electrical & Computer Engineering Department
National Technical University of Athens, Zografou, GR 157 80, Athens, Greece
{gkoutep, maglaris}@netmode.ntua.gr

Abstract. Denial of Service (DoS) attacks are becoming common in the Internet today, employed by malicious Internet users to disrupt or even bring down enterprise networks. Since their first appearances, they have evolved in sophistication, scale, and seriousness of their effects in computer systems and networks. In this paper we examine the main DoS types and their characteristics. We explain why traditional security tools like Intrusion Detection Systems are ineffective and why the problem of countering a Distributed DoS attack is complex, involves various levels of the network, and requires the trust and cooperation between domains. We then look into the solutions offered so far, both practical and research ones. We go through the process of detecting such an attack and lay down a plan for response, manual or automated. Finally, we make a brief review of a system we are currently developing and aims to automate the whole process of attack detection and response. The approach, except for being an alternative solution, highlights the requirements for effective DoS containment.

1. Introduction

Having relatively recently appeared in the security scene, Denial of Service (DoS) and Distributed Denial of Service (DDoS) attacks pose a serious and evolving threat to any networked computer system. Their distinguishing characteristic is that they do not attempt to break into the target computer systems, take control of them or perform information stealing of any kind, like other more "conventional" attacks. Their aim is the disruption of normal operations down to their complete halt. The target is not the system itself but its ability to offer useful services, hence the title of the attack. Targets may range from individual systems to whole domains under attempts to be denied their commercial networking presence. DoS could also be a part of full-scale cyber-warfare confrontations.

DoS attacks fall in two categories: (a) the ones that target a specific system, using certain internal vulnerabilities or trying to overwhelm its processing abilities; another case when a system's vulnerabilities are exploited against it, (b) the ones that target network connectivity on the victim domain. Denial of Service attacks have started as bugs that although could not be exploited for trespassing in systems they were still usable for bringing services down remotely, a malicious alternative to gaining access.

Small ambiguities in the network protocols and their implementations also offered ground for exploitation because when appropriately formed packets reached target systems they could result to their halt. In the evolution of DoS attacks, the network played a crucial role for delivering them in the beginning and by itself becoming the medium that produced and amplified the attack later on.

Although various host attacks were developed and introduced in the late ’90s, it was the utilization of some of the Internet properties that extended the problem and increased its consequences. The malicious users turned to hijacking computers of any size, capabilities and geographical distribution and then using them to stage distributed attacks. They also utilize address spoofing to conceal their origins and the Internet protocol characteristics to amplify their effects. A series of attacks on high profile commercial targets in February of 2000 [1] marked the issue as a serious and threatening problem, able to influence even very powerful systems or high bandwidth networks. The events even prompted a meeting between the US President and members of Internet, e-commerce companies, civil liberties organizations, and security experts to jointly announce actions strengthening Internet and computer network security [2]. More recently, some companies had to completely suspend operations due to continuous interruption to their Internet connectivity [3].

One more characteristic of Denial of Service attacks is that they can affect active networking equipment, like routers. Being specialized computing devices these are usually thought of as “inaccessible” and thus safe. However, they have network connectivity and like ordinary computer systems they include an operating system, many times not free of bugs and vulnerabilities. Although events of unauthorized router access are quite uncommon, the network connection exposes these operational problems to DoS attacks.

In summary, as a security threat DoS attacks present a different paradigm, that of "incapacitating" the victim even without any further goals and have thus to be opposed in non-standard ways; they require new response approaches. Furthermore, good administration and security vigilance, although quite effective practices in other types of attacks, they have proved to be incapable of completely preventing the threat, particularly that coming from DDoS.

2. Host attacks

These comprise the *basic* Denial of Service attack since they usually do not involve more than one attacker and a specific target machine. The first Denial of Service attacks were direct derivatives of methods used for unauthorized access. System bugs or vulnerabilities, which cannot otherwise be used for logging on to the machine or stealing of information may still prevent proper operation or disable a system. The malicious users that are denied access turn to operation disrupting as a way to affect their target.

Some such attacks are:

- *Buffer Overflow Attacks*. They are examples of the dual usage of some penetration methods. Through careful manipulation of the inputs on a poorly written program it is possible to override the machine's defenses by writing directly to the memory stack. The result may be either the establishment of an access point to the machine or the stoppage of normal operations.

- The *Land IP DoS Attack* [4] where a spoofed packet with the SYN flag set is sent to an open port of a host, setting as source the same host and port. This usually causes the machine to halt.

- The *Teardrop Attack* [5] exploits an overlapping IP fragment bug present in various TCP/IP implementations. It sends IP fragments to a network-connected machine causing the TCP/IP fragmentation re-assembly code to improperly handle overlapping IP fragments. The affected machine usually hangs up needing a restart.

The next step in the evolution of these types of attacks was targeting the victims with a multitude of legitimate, but resource consuming requests to be served. The attack is on the victim's recourses (memory, CPU load, etc.) and not on the networking connection, which is being used as the medium for delivering it. Variations of this type of attacks include:

- The *SYN Flooding Attack*, where a lot of connections to a server are left half-open (in the process of TCP three-way handshakes). With each new connection request the server has to commit new resources up to their complete starvation.

- The *Ping Flooding Attack* saturates the target with ICMP echo-request packets that have to be answered. Actually, this may also serve as an attack against the networking connection if the available bandwidth is not enough to handle the flow of packets. A variation of this, the *Smurf Attack* utilizes an intermediate stage, where the ping flow is "amplified" by being first sent to a number of network broadcast addresses with the victim's return address in the packets.

DoS attacks against single hosts are easy to detect, through a Network Intrusion Detection System, checking for the characteristic malicious packet signatures, or on the target itself by a Host Based IDS. Defense against them involves the usual administrator's practices of keeping the systems up to date and following all the latest security developments (see Bugtraq [6]). Filtering of harmful content may occur by setting appropriate filters on the host or, even better on the border router/firewall system. Malicious user detection, however, is much more difficult since it is common practice to use source address spoofing, writing on purpose the wrong sender IP address on the attack packets, thus eluding tracing-back attempts.

3. Network attacks – Distributed DoS

These comprise the basic Denial of Service attack since they usually do not involve more than one attacker and a specific target machine. The first Denial of Service attacks were direct derivatives of methods used for unauthorized access. System bugs or vulnerabilities, which cannot otherwise be used for logging on to the machine or stealing of information may still prevent proper operation or disable a system. The malicious users that are denied access turn to operation disrupting as a way to affect their target. The amplification process in the Smurf Attack and other such events demonstrated the effectiveness of using the distributed Internet infrastructure for producing and delivering attacks. This way malicious packet flows may get magnified so much, that even whole network domain connections can become overwhelmed and unavailable for ordinary traffic. Contrary to what may be though of high-bandwidth connections, some hundred of persistent flows are enough to knock a large network off the Internet. One significant detail of this attack is that incoming traffic has to be controlled, outside the victim's domain, at the upstream providers. Even if the traffic anomaly is detected early the matter of controlling and stopping it falls out of the victim's constituency and direct manipulation abilities. By comparison, a specific machine being targeted can easily be patched, protected at the border router/firewall, or even disconnected from the network.

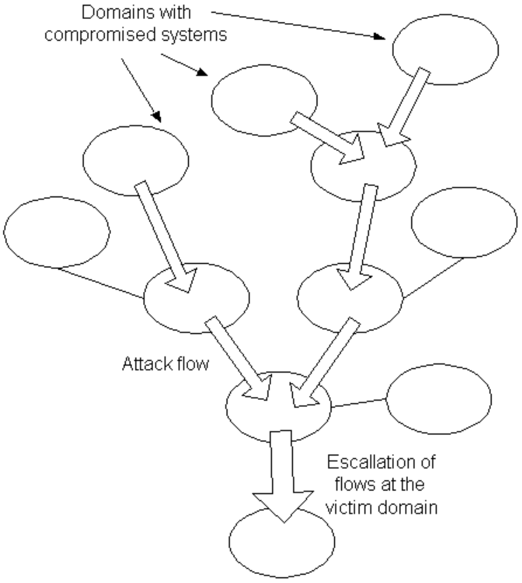

Fig. 1

The DDoS attack is a multi step process with attackers usually exploiting a number of hacked or trojaned computer systems. Initially, through active penetration they install small footprint attack code on a number of machines. These, first stage masters undertake the next phase of establishing an attack infrastructure. They actively scan and identify vulnerable machines, ranging from home computers to big systems and identify vulnerable machines, ranging from home computers to big systems and install there the code that will perform the actual attack. The typical methods of viral or worm infection also apply for spreading these programs. Remotely controlled by the attacker the infrastructure remains inactive for as long as it's necessary and the dormant programs are also referred to as "zombies". When instructed, they activate a flow of packets against the victim network. Although small in scale and difficult to detect near the sources, the flows have a cumulative devastating effect when they reach their target. Fig. 1 summarizes this effect.

The control – command line may have many levels, with the attacker directing a small number of masters and them instructing a greater number of the agents that perform the actual attack. The attacker also has the flexibility to alternate attack targets and traffic sources, may activate and deactivate some of them at will, or adjust the characteristics of the flows according to the sites' reactions , rendering many types of filtering defenses useless.

Examples of such attacks are the ones performed by tools like "Trinoo", "Stacheldraht", and "TFN2K", the so-called *rootkits*. By the latter term we refer to ready-made packages of hacker tools that automate the tasks of attack and code installation to the agents. The particular tools identify and attack vulnerable machines that are then used to further propagate the offensive software. They create a DoS network of hacked machines, controlled remotely by messages from the attacker on specific protocols and ports. In Fig. 2 a multi-tier attack is illustrated. When instructed they launch their flow of packets against the victim. These may be randomized TCP, UDP, and ICMP packets.

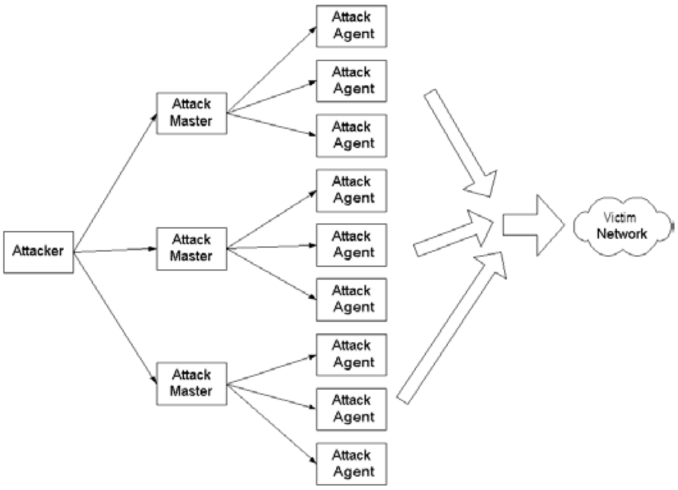

Fig. 2

Except from having many levels of command on the attack network it is also possible to utilize many levels on its delivery; the offensive machines do not have to send the flood of packets themselves but rather construct legitimate requests to Internet servers with the victim's return address. That was initially demonstrated by the *Smurf* attack where ICMP echo requests (with victim's return IP) are sent to broadcast addresses. A newer type "reflects" TCP connection request (SYN) packets on Internet servers (or even routers) that "reply" with TCP SYN-ACK packets directed to the victim. An example is presented in Fig. 3.

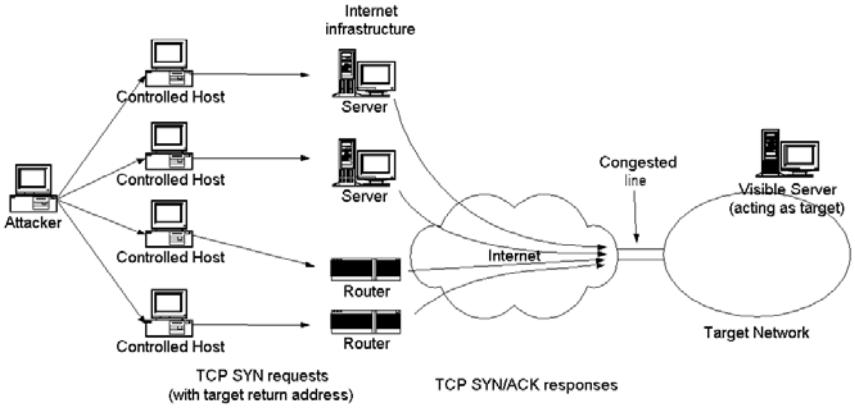

Fig. 3

DDoS attacks present an interesting management challenge since their nature makes them difficult to stop by the efforts of a single site. Factors that contribute to this are (a) the practice of attackers to spoof packet source IPs, (b) the possibility of the attack initiating from a wide range of networks worldwide, and (c) the inability of a domain to enforce incoming traffic shaping; detected malicious flows can be blocked locally but the assistance of the upstream network is still needed in order to free the resources occupied by them on the incoming link.

4. Detection, Prevention, and Reaction

Detecting a Denial of Service attack is just one of the actions in the process of countering it. If it is a single host – single target attack then the detection relies on the ability of the Intrusion Detection system deployed and protection depends on how effective the Firewall system will be in blocking the particular malicious packets. Since these types of attacks depend upon new system vulnerabilities that evolve constantly it is not always certain that the IDS and Firewall systems will be up to date with their rules. So a number of standard administration procedures must also be in place to ensure individual system security or even full recovery in the event of an attack succeeding in bringing a mission critical system down. Some effective actions should be:

- Maintain a backup of the system, or, even better, a standby machine to take over the role of the affected one. Ideally the two will be based on different OSes and make the switch automatically should the primary become unavailable.

- Standard good administrative practices also apply in this case. The administrator should monitor all corresponding security mailing lists and announcements and install patches and fixes as soon as they become available for his architecture.

- Border defenses, Intrusion Detection and Firewall systems must also be kept up to date to detect and block attacks. A high number of IDS alerts, even if unconnected with an attack, may serve as a warning of potentially malicious intentions and data collection on prospective targets.

Reaction to Distributed Denial of Service that targets the networking infrastructure is a complex procedure. As explained earlier, any effective response requires cooperation between sites. The procedure takes place as follows:

- The specific characteristics of the malicious flows have to be determined at any point that the attack traverses. This will enable the installation filters exactly suited to the attack. Determining attack characteristics are the protocols used, the ports, and very rare (since they are usually spoofed) source addresses. These comprise and characterize the traffic *aggregate* of the attack (as described in [9]).

- The attack characteristics have then to be communicated to the network(s) on the attack path. Since the network connection may be completely out this has usually to be done manually and is an uncertain and time-consuming procedure.

- The effectiveness and success of this process depends heavily on the upstream network administrator's availability and good will, as well as the service policies there. According to the site's security policies the actions that will be implemented usually consist of setting up tailor-made blocking or throttling filters on active network components.

- The filtering process requires contact between the victim and the upstream network to check the effectiveness of the procedure. The implemented filters require constant monitoring and adjusting to shifting attack patterns. Finally they have to be deactivated in the end of the event, especially if they hinder normal traffic patterns, because the attack makes use of them.

Unfortunately, no matter how effective this response will be, the bandwidth penalty is still present throughout all the affected networks. To alleviate the resulting congestion extra steps must be taken and contacts be made between the sites on the attack path, to further resolve the situation. Obviously the further we move from the victim, the more dispersed this procedure becomes and there is less immediate interest from the domains to help.

The prevention and reaction measures that should be taken at a network are:

- Ensure that an attack will be interpreted right (not as a network outage) and as soon as it appears. The type and traffic characteristics, target, and origin of the attack have also to be determined as soon as possible. So, it helps to have security aware management personnel, capable and available to react to an event 24x7.

- Have a prepared action plan for the case of a DoS/DDoS attack, including emergency phone lines and possibly out-of-band (or dialup) small-bandwidth connections.

- Configure your router to do egress filtering, preventing spoofed traffic from exiting your network. [7] has more information on this.

- Have established contact points with upstream networks. Ensure that the provider's policies provision for actions in that case of the need for active response.

- Furthermore, it helps to have contacts with CERT organizations that may undertake the task of further propagating the reaction process to more networks on the attack path, nationally or internationally.

In summary, the requirements for an effective response to a DDoS attack are: (a) Early detection both at the victim site and at upstream stages, (b) flow of incident information between domains, effective and timely domain cooperation but according to each domain's policies (c) quick, automatic, and effective response in as many domains on the attack path as possible, and (d) avoiding extra network overloading due to these communications [10].

5. Research directions

As a research issue, countering DoS attacks is also divided in the host-based and network based problems Single-host attacks constitute challenges on intrusion detection and response not much different than any other security threat. Actually the method of attack delivery, single packets packets, directed to specific machines and services are easier to detect that many instances of premeditated, gradual and persistent to break-in attempts. Once an attack's characteristics have been established it's rather easy to renew signatures of a host or network IDS to detect it . New and unknown types of attacks can be determined by anomaly patterns on the characteristics of traffic, network connections or host system calls. The last method, of determining anomalies on the line of system calls mimics the human immune system and can determine if a certain sequence of calls has higher or lower anomaly factor [11], [12]. Windows of time and number of observed calls can fine-tune the procedure towards very accurate results either for DoS or break-in attempts. Generally the "Anomaly" detection methods cannot offer remedy for attacks that can be executed with a single packet that some times can have destructive results, whether it is detected or not.

Distributed DoS against networking resources presents a bigger and more complicated challenge since significant problems exist in the areas of (a) detection, (b) source tracing, and (c) traffic flow control and attack suppression. Most importantly, as presented earlier, a single domain cannot undertake the task of stopping a DDoS.

Though capable of attack identification and early warning, conventional Intrusion Detection (ID) Systems usually cannot offer domain traversing and scalable response capabilities. One extra weakness is that, although there are solutions for message exchange between them (such as IETF's Intrusion Detection Working Group [13]), they lack the underlying cooperation framework to transmit specific requests for a suitable response on another domain. Even if this would be possible, the security concerns would deny any direct manipulation of the necessary networking equipment. Furthermore, the IDS would have to know the remote devices' topology and interfaces in the other domain to issue appropriate commands.

On single, big-size ISPs a number of approaches have been proposed that may offer some remedy to DDoS Attacks but all have some negative consequences:

1. Since the target of the attack can easily be identified (or just be communicated) it is possible to stop all its traffic through the ISP (route all its traffic from a central point to a "dead end"). The positive result is that generally bandwidth consumption at the ISP level is alleviated. Additionally hop-by-hop tracking can be performed in the rest of the affected network. The disadvantages of this solution are that (a) the victim does not have any improvement in its condition since all the traffic (even legitimate) towards this site stops (inconsiderate of "good", "bad", or "poor" nature, as defined in [9]) and (b) there is no effort to trace the attack back to its sources, react at the ISP's border routers or on any upstream networks.

2. In ref [14] there has been a proposal for setting up an overlay network called *CenterTrack* for tracking the DoS floods. Tunnels are set up from the domain's edge routers to a central point in the network. Once a victim has been established all the traffic to him is still delivered but routed through the overlay network, achieving hop-by-hop trace-back to those edge routers that are transcended by the attack traffic. The main advantage of the method is that specialized diagnostic features are required only at the edge routers. The establishment of the traffic tunnels, however, consumes resources, increases the management complexity on critical systems and has high overhead.

3. In [8] J. Ioannidis, et al. suggest a solution of controlling the high bandwidth traffic flows (called "aggregates") that comprise a DDoS attack. Once they have established the characteristics of the DDoS the malicious traffic they move on to block them at the routers using tailored filters. They then communicate their findings hop-by-hop, between cooperating routers using the special *Pushback* protocol. The latter is an effort towards the standardization of such methods. As a result they follow the malicious aggregates closer to their source and control their bandwidth allocation in many points. The main disadvantage of the method is that it has to consume processing resources on all the routers towards the source. To achieve the shift to the next hop all the routers must employ the same technique.

4. Other approaches attempt to measure changes in traffic flow patterns on edge routers. One such open-source tool, called *Panoptis* [15] correlates recent packet and flow counts to derive which of the edge router interfaces are involved in the attack. The measurements are supplied from the Cisco router Netflow instrumentation. Specific filters can then be set up manually on the specific interfaces. In this type of solution we have good scaling factor but a big amount of historical data must be kept and all processing must take place in locally to avoid bandwidth consumption by accounting data.

Our team attempts to fulfill all the requirements for countering DDoS attacks by following another approach with the introduction of a framework for effective cooperation between whole domains. We propose the *Cooperative IDS Entity*, a software system (installed on each domain) that constitutes a trusted peering point within the proposed distributed framework. The IDS Entities are deployed on top of each local IDS hierarchy receiving messages from it and exchanging messages with their peers on other domains, primarily by multicast messages. Special measures are taken to ensure the integrity and security of these exchanges. Furthermore each Entity has a limited local response capability. The inference engine of the IDS Entity combines local and remote information about on-going security events and responds to them locally according to administrator-defined policies. Our work focuses on a higher level approach to the problem. We try to provide a cooperation-enabling infrastructure for domains and we are independent from the Intrusion Detection part, although we are coupling our prototype with a router DoS detection tool to provide attack identification. Intrusion Detection is a matter open to the choices of the individual domains [17].

References

1. "ALERT 00-034", National Infrastructure Protection Center, February 10, 2000, http://www.nipc.gov/warnings/alerts/2000/00-034.htm

2. "Strengthening Cyber Security through Public-Private Partnership", Center for Democracy and Technology, February 15, 2000, http://www.cdt.org/security/dos/000215factsheet1.shtml

3. "Cloud Nine blown away, blames hack attack", The Register, January 22, 2002, http://www.theregister.co.uk/content/6/23770.html

4. "The LAND attack (IP DOS)", Insecure.org, 20 November 1997, http://www.insecure.org/sploits/land.ip.DOS.html

5. "CERT Advisory CA-1997-28 IP Denial-of-Service Attacks", CERT Coordination Center, December 16, 1997, http://www.cert.org/advisories/CA-1997-28.html

6. Bugtraq Mailing List http://www.securityfocus.com/

7. "Egress Filtering v 0.2", Global Incident Analysis Center, Special Notice, SANS Institute, http://www.sans.org/y2k/egress.htm

8. J. Ioannidis and S. Bellovin, "Pushback: Router-Based Defense Against DDoS Attacks", NDSS, February 2002.

9. R. Mahajan, S. Bellovin, S. Floyd, J. Ioannidis, V. Paxson, and S. Shenker, "Controlling High Bandwidth Aggregates in the Network", draft, February 2001.

10. G. Koutepas, F. Stamatelopoulos, and B. Maglaris, "Efficiency and Performance Issues in Distributed Intrusion Detection Systems", Applied Telecommunication Symposium 2002 (ATS 02), San Diego, CA, USA, April 2002.

11. Stephanie Forrest, Styeven A. Hofmeyer, Amil Somayaji, “Computer Immunology”, Communications of the ACM, October 1997.

12. P. Astithas, V. Pappas, B. Maglaris, "Detecting Intrusions by Monitoring System Processes", in Proceedings of the 8th HPOVUA Plenary Workshop on Network and Systems Management, Berlin, Germany, June 2001.

13. Intrusion Detection Working Group (IDWG) of the IETF, http://www.ietf.org/html.charters/idwg-charter.html

14. R Stone, “CenterTrack: An IP Overlay Network for Tracking DoS Floods”, 9th USENIX Security Symposium, Denver Col., USA, August 2000.

15. C. Kotsokalis, D.Kalogeras, and B. Maglaris, "Router-Based Detection of DoS and DDoS Attacks", HP OpenView University association (HPOVUA) Conference '01, Berlin, Germany, June 2001.

16. Panoptis home page. http://panoptis.sourceforge.net/

17. G. Koutepas, F. Stamatelopoulos, B. Maglaris “A Trans-Domain Framework Against Denial of Service Attacks: Communication and Cooperation Aspects for Efficiency and Effectiveness”, Submitted to the 10th Annual Network and Distributed System Security Symposium, San Diego, California, February 2003.

Georgios Koutepas

Mr. Koutepas is a PhD. candidate at the National Technical University of Athens, Greece. He holds a BA. from the NTUA. He is currently doing research on effective trans-domain DDoS detection and containment. He can be reached at gkoutep@netmode.ntua.gr.

Basil Maglaris

Dr. B. Maglaris is professor at the dept. of Electrical and Computer engineering at the National Technical Univerisity of Athens, Greece. He holds a Ph.D. degree in Electrical Engineering & Computer Science from Columbia University, New York. Professor Maglaris is teaching and performing research on communication networks, with focus on operations and management of communication networks, telematics applications, queuing theory, performance models of computer network protocols, and optimal design and planning algorithms of communication networks. He teaches courses on queuing theory, management and control of telecommunications networks, and tele-traffic source models at the Department of Electrical & Computer Engineering of NTUA. He can be reached at maglaris@mail.ntua.gr.