Defenses against denial-of-service Attacks There are a number of steps that can be taken both to limit the consequences of being the target of a DoS attack and to limit the chance of your systems being com promised and then used to launch DoS attacks.In general, there are four lines of defense against DDoS attacks [PENG07, CHAN02]: o Attack prevention and preemption (before the attack): These mechanisms enable the victim to endure attack attempts without denying service to legiti mate clients.RFC 2827, Network Ingress Filtering: Defeating Denial-of-service attacks which employ IP Source Address Spoofing, 8 directly makes this recommendation, as do SANS, CERT, and many other organizations concerned with network security.The provision of significant excess network bandwidth and replicated distributed servers is the usual response, particularly when the overload is anticipated.Similarly, when popular sporting events like the Olympics or Soccer World Cup matches occur, sites reporting on them experience very high traffic levels.In addition, preven tion mechanisms modify systems and protocols on the Internet to reduce the possibility of DDoS attacks.Hence one of the fundamental, and longest standing, recommendations for defense against these attacks is to limit the ability of systems to send packets with spoofed source addresses.This type of filtering can be implemented using explicit access control rules in a router to ensure that the source address on any customer packet is one allocated to the ISP.In particular, if an attacker can direct a large enough volume of legitimate traffic to your system, then there is a high chance this will overwhelm your system's network connection, and thus limit legitimate traffic requests from other users.Classically, a posting to the well-known Slashdot news aggregation site often results in overload of the referenced server system.These either obscure the originating system of direct and distributed DoS attacks or are used to direct reflected or amplified traffic to the target system.This filtering needs to be done as close to the source as possible, by routers or gateways knowing the valid address ranges of incoming packets.An ISP knows which addresses are allocated to all its customers and hence is best placed to ensure that valid source addresses are used in all packets from its customers.7.6 / Defenses against denial-of-service Attacks 261 Alternatively, filters may be used to ensure that the path back to the claimed source address is the one being used by the current packet.Techniques include enforcing policies for resource consumption and providing backup resources available on demand.True ingress filtering rejects outside packets using source addresses that belong to the local network.There is very little that can be done to prevent this type of either 260 Chapter 7 / Denial-of-Service Attacks accidental or deliberate overload without also compromising network performance.However, this method typically does not yield results fast enough, if at all, to mitigate an ongoing attack.A critical component of many DoS attacks is the use of spoofed source addresses.8 Note that while the title uses the term Ingress Filtering, the RFC actually describes Egress Filtering, with the behavior we discuss.o Attack detection and filtering (during the attack): These mechanisms attempt to detect the attack as it begins and respond immediately.o Attack source traceback and identification (during and after the attack): This is an attempt to identify the source of the attack as a first step in preventing future attacks.Typically this is the ISP providing the network connection for an organization or home user.This is regularly done for popular sporting sites.However, this response does have a sig nificant implementation cost.Detection involves looking for suspicious patterns of behavior.Response involves filtering out packets likely to be part of the attack.It is important

to recognize that these attacks cannot be prevented entirely.This has led to the terms slashdotted, flash crowd, or flash event being used to describe such occurrences.o Attack reaction (after the attack): This is an attempt to eliminate or curtail the effects of an attack.This provides protection against only a small number of attacks.Indeed, this sometimes occurs by accident as a result of high publicity about a specific site.This minimizes the impact of the attack on the target.We discuss the first of these lines of defense in this section and consider the remaining three in Section 7.7.