

Note: This is an historic document. We are no longer maintaining the content, but it may have value for research purposes. Pages linked to from the document may no longer be available.

Results of the Distributed-Systems Intruder Tools Workshop

Pittsburgh, Pennsylvania USA

November 2-4, 1999

Published at the
CERT® Coordination Center
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213

December 7, 1999

CERT and CERT Coordination Center are registered in the U.S. Patent & Trademark office by Carnegie Mellon University.

Contributors

The ideas in this paper were jointly developed by participants in the Distributed-Systems Intruder Tools Workshop. Their intellectual contributions and their spirit of cooperation made the workshop a success. Among the many participants who contributed to this paper are the following:

Jon David
AT&T Information Security Center

Cory Cohen, Kathy Fithen,
Kevin Houle, Tom Longstaff,
John McHugh, Eric Mitchell,
Rich Pethia, Jed Pickel,
Tim Shimeall, Dara Sewell (resident affiliate)
CERT® Coordination Center

Bradley Frank
Ken Rowe
Cisco Systems, Inc.

Brian Dunphy
Sean McAllister
DoD CERT

Sammy Miguez
Infrastructure Defense

Pat Becker
Internet Security Systems, Inc.

Sven Dietrich
Aghadi Shraim
NASA Goddard Space Flight Center

John Green
NSWC (Naval Surface Warfare Center) SHADOW Team

Richard Forno
Network Solutions, Inc.

Kenneth R. van Wyk
Para-Protect®, Inc.

Kathleen Kimball
George Weaver
Penn State University

Clarissa Cook
Robert Stone
UUNET

Richard A. Kemmerer
University of California, Santa Barbara

David Dittrich
University of Washington

N.L.

Results of the Distributed-Systems Intruder Tools Workshop

Pittsburgh, Pennsylvania USA
November 2-4, 1999

Executive Summary

On November 2-4, 1999, the CERT® Coordination Center invited 30 experts from around the world to address a category of network attack tools that use distributed systems. Several tools are in use now, and the technology is maturing. As a result, a single, simple command from an attacker could result in tens of thousands of concurrent attacks on one or a set of targets. The attacker can use unprotected Internet nodes around the world to coordinate the attacks. Each attacking node has limited information on who is initiating the attack and from where; and no node need have a list of all attacking systems. Damaged systems include those used in the attack as well as the targeted victim. For the victim, the impact can be extensive. For example, in a denial-of-service attack using distributed technology, the attacked system observes simultaneous attacks from all the nodes at once – flooding the network normally used to communicate and trace the attacks and preventing any legitimate traffic from traversing the network.

Distributed intruder technology is not entirely new; however, it is maturing to the point that even unsophisticated intruders could do serious damage. The Distributed-Systems Intruder Tools (DSIT) Workshop provided a venue for experts around the world to share experiences, gain a common understanding, and creatively brainstorm possible responses and solutions *before* the dissemination of the maturing attack tools – and attacks themselves – become widespread.

One consideration is the approach typically taken by the intruder community. There is (loosely) organized development in the intruder community, with only a few months elapsing between “beta” software and active use in attacks. Moreover, intruders take an open-source approach to development. One can draw parallels with open system development: there are many developers and a large, reusable code base. Intruder tools become increasingly sophisticated and also become increasingly user friendly and widely available. As a result, even unsophisticated intruders can use them.

There has already been some public discussion in the intruder community about distributed attack tools while development continues. In their development, intruders are using currently available technology to develop new technology. For example, they are building on previous scanning technology and automated intrusion tools to create more

powerful intrusion tools. One concern of workshop participants is that in a relatively short time, it may be possible for unsophisticated intruders to gain control of and use systems distributed across significant portions of the Internet for their attacks.

This paper is one outcome of the DSIT Workshop. In it, workshop participants examine the use of distributed-system intruder tools and note that current experiences have highlighted the need for better forensic techniques and training, the importance of close cooperation, and a concern for the rapid evolution of intruder tools. They provide information about protecting systems from attack by the tools, detecting the use of the tools, and responding to attacks. The paper includes suggestions for specific groups in the Internet community:

- managers
- system administrators
- Internet service providers (ISPs)
- incident response teams (IRTs)

The suggestions address actions each group should take immediately, along with actions for the short term and long term. They also remind readers that the security of any network on the Internet depends on the security of every other network. The widely varying implementation of security measures is what often makes a distributed attack successful.

The workshop participants hope that the information offered here will help reduce the impact of distributed attack tools on the Internet as those tools mature.

Results of the Distributed-Systems Intruder Tools Workshop

1. Introduction

On November 2-4, 1999, the CERT® Coordination Center (CERT/CC) invited 30 experts from around the world to address a category of network attack tools that use distributed systems in increasingly sophisticated ways. Intruders are maturing an attack technology that goes beyond using individual systems as the starting point for an attack. Rather, they can potentially use tens of thousands of unprotected Internet nodes together in order to coordinate an attack against selected targets. Each attacking node has limited information on who is initiating the attack and from where; and no node need have a list of all attacking systems. For the victim, the impact can be extensive. For example, in a denial-of-service attack using distributed technology, the attacked system observes simultaneous attacks from all the nodes at once – flooding the network normally used to communicate and trace the attacks and preventing any legitimate traffic from traversing the network.

Distributed intruder technology is not entirely new; however, it is maturing to the point that even unsophisticated intruders could do serious damage. In the past, intruders have used IRC robots to control remotely networks of compromised machines. In addition, *fapi*, a denial-of-service (DoS) tool that appeared early in 1998, works in a similar way to some of the tools we are now seeing, but it was not as sophisticated or as widely used.

During the Distributed-Systems Intruder Tools (DSIT) Workshop, participants discussed a large number of approaches to preventing, detecting, and responding to distributed-systems attacks. The CERT/CC specifically invited technical personnel that could contribute technically to the solutions regardless of their position in their home organization or political stature in the community. Thus, the workshop effectively provided a venue for experts around the world to share experiences, gain a common understanding, and creatively brainstorm possible responses and solutions to this category of attack *before* the dissemination of the attack tools – and the attacks themselves – become widespread.

One consideration is the approach typically taken by the intruder community. There is (loosely) organized development in the intruder community, with only a few months elapsing between “beta” software and active use in attacks. Intruders are actively developing distributed tools to use the many resources on the network; this has become easier because of the large number of machines “available for public use” – that is, vulnerable to compromise and, thus, available for use by anyone who can exploit the vulnerabilities. Moreover, intruders typically take an open-source approach to development. One can draw parallels with open system development: there are many developers and a large, reusable code base. Intruder tools become increasingly sophisticated and also become increasingly user friendly and widely available. As a result, even unsophisticated intruders can use the available tools to identify and take advantage of a large number of vulnerable machines.

There has already been some public discussion in the intruder community about the distributed attack tools while development continues. Intruders are using currently available technology to develop new technology. For example, they are building on previous scanning technology and automated intrusion tools to create more powerful intrusion tools. One concern of workshop participants is that in a relatively short time, it may be possible for unsophisticated intruders to gain control of and use systems distributed across significant portions of the Internet for their attacks.

As noted in the letter of invitation to the participants,

So far, we have seen only limited use of these new tools, but we believe it won't be long before the tools will move from the development by sophisticated intruders into wide use by the large population of less sophisticated intruders. When this happens, all of us will face new issues with impact on security, incident response, and future technology. ...

I believe that security experts need to act now, before the tools are in widespread use. During the workshop, we hope to analyze these new attack tools; explore their possible evolution and kinds of impact we might see from their use; and outline techniques that can be used to detect, respond to, and recover from attacks.

One strong response to the workshop from the participants is that prior to the workshop, there was no way for the technical staff at important critical infrastructure sites to communicate the threat to management. The participants could understand the problem from an isolated perspective, but it was not until the workshop brought them together that the true nature of the threat was understood and could then be communicated to the management at their home organizations. In many cases, the resulting briefs given to the home organization (including government agencies, critical commercial providers, and university researchers) provided the first and best view of the nature of the changing threat in using networked systems. Finally, this paper, which summarizes output from the workshop, enables the Internet community to gain similar understanding and to take action.

In the next section, workshop participants examine the use of distributed-system intruder tools. Later sections provide information for specific groups in the Internet community:

- managers
- system administrators
- Internet service providers (ISPs)
- incident response teams (IRTs)

The workshop participants hope that the information offered here will help reduce the impact of the attack tools on the Internet as those tools mature.

2. Recent Activity Involving Distributed Attack Systems

Distributed systems based on the client/server model have become increasingly common. In recent months, we have seen an increase in the development and use of distributed sniffers, scanners, and denial-of-service tools. Attacks using these tools can involve a large number of sites simultaneously and be focused to attack one or more victim hosts or networks.

During the second half of 1999, several sites reported denial-of-service attacks involving distributed intruder tools. While some of the details presented here are specific to the incidents that were observed, the overall distributed strategy can be applied to attacks other than denial of service. The description in this section concentrates on the distributed aspects of the incidents while omitting unnecessary details.

As shown in the figure below, in a typical distributed attack system, the “intruder” controls a small number of “masters,” which in turn control a large number of “daemons.” These daemons can be used to launch packet flooding or other attacks against “victims” targeted by the intruder.




Figure 1 – Distributed-Systems Attack

In the incidents that have occurred so far, daemons were installed at several hundred sites, typically through the exploitation of well-known vulnerabilities that lead to root privileges on the compromised machines. Though some implementations of the daemon program do not require root privileges to launch attacks, in practice most of the daemons were concealed by the installation of “root kits” designed to hide evidence of the intrusion. Intruders have also sometimes used system facilities such as “cron” to ensure that a daemon would continue to run even if one instance of it were deleted or the system was rebooted.

There are indications that the processes for discovering vulnerable sites, compromising them, installing daemons, and concealing the intrusion are largely automated, with each step being performed in “batch” mode against many machines in one “session.” Daemons have been discovered on a variety of operating systems with varying levels of security and system management.

Once installed and operating, the daemon announces its presence to several (usually three or four) predefined masters and awaits further commands. The master program records that the daemon is ready to receive commands in an internal list, which can be retrieved by the intruder. Lists recovered from incidents have included hosts in several different nations. Masters can cause daemons in the list to launch attacks, shut down gracefully, or even announce themselves to a new master server. Intruders have used cryptographic techniques to conceal the information recorded by the master daemons.

Upon command from an intruder, the master can issue attack requests to the daemons in its list. These requests contain information about the requested attack, such as the address of the victim, the duration, and other parameters. Upon receipt of the request, the daemon proceeds to attack the victim, usually by flooding the victim with packets. No further contact from the master is necessary.

The master programs frequently operate as ordinary user programs on compromised hosts, where their activity can easily be hidden. Unlike the daemon programs, which are intended to be run on sites with a substantial network capacity, traffic to and from the master program is limited to control messages.

In one incident reported to the CERT Coordination Center, a flooding attack was aimed at a major university. This attack involved several hundred daemons scattered over a wide variety of locations, and it generated enough traffic to disable the university’s Internet connectivity for a period of several days.

Several incidents have indicated that intruders are actively seeking systems with good network connectivity for compromise and installation of the daemon program. The indiscriminate installation of daemons on any system with a significant network capacity has included systems whose compromise could have life-threatening consequences.

The experiences of those who reported early attacks highlight the need for better forensic techniques and training, the importance of close cooperation, and concern for the rapid evolution of intruder tools.

- **Better forensic techniques and training** – Detecting and eliminating master programs is a critical part of disabling a distributed intruder system, but unfortunately the masters often do not leave obvious signs of intrusion on the system where they are installed. In most cases, the master hosts were identified after forensic examination of daemons involved in a denial-of-service attack. This forensic analysis was expensive and limited to a few knowledgeable people with experience in the field, but ultimately most of what we know today about how the systems work is a result of this analysis. Forensic techniques and training must be available to a much larger audience to respond to these attacks in the future.
- **Close cooperation and communication** – Prior to the workshop, many participants had incomplete information regarding the tools and methods used by intruders in this kind of attack. By sharing their knowledge, they were able to establish a more complete understanding of distributed intruder tools.
- **Rapid evolution of intruder tools** – The intruder tools encountered in the incidents leading to the creation of this document changed substantially during the planning of the workshop and have continued to evolve since then. As intruders learn to use established technologies to their advantage, the incident response community needs to be better prepared to meet this challenge.

3. Audience-Specific Information

Managers

For management, the issues related to the ongoing development of distributed attack tools, such as trinoo and tribe flood network (for details, see CERT/CC incident note IN-99-07: http://www.cert.org/incident_notes/IN-99-07.html), center largely around the need to understand fully the ramifications of the intruder tools and to perform impact and organizational risk assessments on a priority basis. The results of these assessments then need to be incorporated into plans such as those for operational guidance, equipment acquisition, service contracts, and equipment configuration.

Planning and coordination before an attack are critical to ensuring adequate response when the attack is in progress. Since the attack methodology is complex and there is no single-point solution or “silver bullet,” resolution and restoration of your systems may be time-consuming. The bottom line for management is that your systems may be subject at any time to distributed attacks that are extremely difficult to trace or defend against.

Although an organization may be able to harden its own systems to help prevent implantation of the daemon portion of a distributed attack tool, there is essentially nothing a site can do with currently available technology to prevent becoming a victim of, for

example, a coordinated network flood. The impact upon your site and operations is dictated by the (in)security of other sites and the ability of a remote attacker to implant the tools and subsequently to control and direct multiple systems worldwide to launch an attack. The result may be reduced or absent network connectivity to your enterprise for extended periods of time, possibly days or even weeks depending upon the number of sites attacking and the number of possible attack networks that could be activated in parallel or sequentially. Therefore, to minimize the effect on business operations, it is important to know and document *in advance* the actions the enterprise will take and the primary contingency contacts who must be notified.

Below are some recommend actions for coping with the potential for an attack using distributed-system intruder tools:

- Become fully informed with regard to the nature of the attacks and the potential ramifications. Senior management should receive direct briefings from security staff in an effort to facilitate full understanding.
- Be cognizant of your own site's security posture. If your site is capable of being easily compromised due to inattention to security issues and your systems are used as either master(s) or daemon(s) for such an attack, it is possible you may share liability for damage caused to victim sites. (Consult with your organization's legal advisors and inform them of the attacks.) The reputation of your enterprise may also be at risk from the adverse publicity that may result.
- Assess the services that are mission critical for your particular business. Determine the impact upon mission-critical services if Internet connectivity is unavailable for an extended period. Develop contingency plans for continuity of operations in the event of an extended Internet outage. Consider and plan to insure against possible revenue loss due either to lost opportunity (for example, the absence of connectivity to your site for staff members, external customers, and business partners) and in lost sales (for example, an electronic commerce site is flooded and orders cannot be received). Read insurance policies carefully, and seek legal opinion on coverage for distributed-systems attacks.
- Develop an augmentation strategy to provide staff and other resources in the event of an attack. Determine which staff may be needed and where they should report. Be sure there are phone or alternative communications since electronic communication may be difficult or impossible.
- Be sure your staff have the time and resources needed to perform traffic analysis, intrusion detection, coordination with upstream providers, and other activities described under "System Administrators" below.
- Ensure privacy issues associated with log retention and review have been addressed in policy and that adequate analytical information is readily available to critical staff in the event an attack occurs.

- Examine your current policy requirements. In particular, ensure responsibility is defined for 1) enforcing minimum security standards; 2) cutting off users (even executive-level users) whose accounts may have been compromised or are at risk; and 3) disconnecting uncontrolled Internet connections.
- Be sure that all levels of management understand and are held accountable for security planning and implementation. Be sure that an adequate and enforceable acceptable use policy exists enterprise-wide.
- Realize that the escalating Internet threat environment must be matched by corresponding investments in security. Define security resources in the budget.
- Examine your current network and security architecture. Many sites have optimized connectivity for speed of access, making decisions that complicate security measures. In the escalating threat environment, speed and reliability can be denied unless security is included in the architecture.
- Aggressively develop cooperative relationships to support security across organizations and policy to govern those relationships. To deal effectively with distributed agents, your organization may need to cooperatively support security at other Internet sites. Internet service providers and incident response organizations should be supported.
- Pressure vendors to provide more security in their default services and configurations. Simply correcting known vulnerabilities in new releases would reduce the population of candidate sites for intruders. Ask your vendors specifically if they support the capabilities listed in the “Internet Service Providers” section.

Finally, managers need to consider these trends:

- The intruder community is actively developing distributed technology.
- There are multiple categories of existing distributed-systems tools, including distributed sniffers, denial of service, and information gathering.
- In a relatively short amount of time, unsophisticated intruders can acquire sophisticated tools, enabling them to control and use significant portions of the Internet for their attacks.

System Administrators

With the increased sophistication of intruder tools comes the critical need for action. The following table lists actions identified at the Distributed-System Intruder Tools Workshop, along with a suggested time frame for dealing with attacks using distributed-system tools.

	Immediately (< 30 days)	Near Term (30 – 180 days)	Long Term (> 6 months)
Protect	<ul style="list-style-type: none"> ▪ Apply anti-spoofing rules at the network boundary. (This makes your site a less appealing target for intruders.) ▪ Keep systems up to date on patches. ▪ Follow CERT/CC & SANS best practices. ▪ Review boundary security policy to ensure outbound packets are restricted appropriately. 	<ul style="list-style-type: none"> ▪ Establish reference systems using cryptographic checksum tools such as Tripwire®. ▪ Scan your network periodically for systems with well-known vulnerabilities & correct problems that you find. ▪ Evaluate & (possibly) deploy an intrusion detection system (IDS). 	<ul style="list-style-type: none"> ▪ Identify a system administrator with responsibility for each system, who has the authority, training & resources to secure the system. ▪ Deploy resources for host-based intrusion detection. ▪ Provide security training for users. ▪ If you do not have sufficient resources or support to effectively protect systems, lobby for them.
Detect	<ul style="list-style-type: none"> ▪ Look for evidence of intrusions in logs, etc. ▪ Look for distributed tool footprints as described in documents from the CERT/CC or your incident response team. ▪ Enable detection of unsolicited ICMP echo replies & unusually high traffic levels. 	<ul style="list-style-type: none"> ▪ Periodically compare systems to your reference system using cryptographic checksum tools such as Tripwire®. ▪ Run host-based software to detect vulnerabilities & intrusions. 	<ul style="list-style-type: none"> ▪ Develop a system for profiling traffic flows & detecting anomalies, suitable for real-time detection & prevention. ▪ Create and practice a response plan.
React	<ul style="list-style-type: none"> ▪ Report to a predefined list of contacts, approved by management. ▪ Establish detailed, written, management-approved plans for communicating with IRTs, ISPs, & law enforcement. Include out-of-band contacts. ▪ Obtain training & experience in forensic techniques required to analyze compromised systems & identify other hosts involved, such as the master hosts in a distributed network. 	<ul style="list-style-type: none"> ▪ Ensure ability to capture, analyze, & collect forensic evidence accurately & quickly by developing a “forensic toolkit” of tools & programs to assist in forensic analysis. ▪ Work with your ISP to establish a good business relationship, with service-level agreements that identify the ISP’s responsibilities in tracking & blocking traffic during DoS attacks. 	<ul style="list-style-type: none"> ▪ Work with management to ensure that policies are in place that allow appropriate measures against suspect systems. ▪ Work with your ISP to implement improved security requirements & capabilities in your service-level agreement.

Table 1 – Suggestions for System Administrators

Additional comments for system administrators:

When you set up intrusion detection software, ensure that it is both fault tolerant and capable of maintaining logs on a highly saturated network. The definition of a highly saturated network varies from organization to organization. A good metric is the amount of traffic seen divided by the maximum bandwidth available to the organization. Expect to see near 100% capacity during a distributed denial-of-service attack.

In setting up logs, have the ability to parse log information at a high rate. Workshop participants recommend attention be paid to searching based on host name/IP number.

Be able to search at least packet headers for attack signatures.

Finally, look to an incident response team for techniques and information for dealing with distributed attacks and the evolving attack tools.

Internet Service Providers (Network Operators)

For the purposes of this report, an Internet service provider (ISP) is considered to be an entity that operates an Internet backbone that is used to carry traffic between two or more other Internet-connected networks. The term *ISP* refers to commercial network operators, research and education networks, government-operated networks, etc.

The transport and access portions of networks characterize the unique role of an ISP in the context of a distributed-system attack. Packets generated from multiple sources during a distributed denial-of-service attack, for example, are likely to be transported across one or more ISP network backbones en route to the victim site. The access portions of an ISP's network (physical connection points of downstream hosts and networks) may be either components of an attack or the end victim.

Considering only the transport and access portions of ISP networks, a network operator's role in a distributed attack is essentially composed of two things:

1. Identifying and controlling traffic flows from the point the traffic enters the network (ingress) to the point the traffic leaves the network (egress).
2. Ingress filtering at the network edge and/or network borders to prevent origination of packets with spoofed source IP addresses.

In addition to the unique characteristics of the ISP networks, the networked computer systems used by ISPs to deliver services such as DNS, email, and web hosting may be attractive locations for intruders to install distributed-system tools for several reasons:

- Active traffic patterns may obscure the use of attack tools.
- Close proximity to high-capacity network backbones enables attacks to have a high impact.

ISP systems themselves may also be high-impact targets for distributed-system attacks. People and systems depending on an ISP's services tend to use shared resources at some level. A carefully targeted attack on one or more critical shared resources may affect a large number of Internet users.

The issues facing an ISP with regard to its networked computer systems being used in an attack, or being the target of an attack, are otherwise not unique and can be considered to be on par with issues faced by system and network administrators at other Internet sites (see the section for system administrators).

During an ongoing attack, an ISP may need to trace traffic flows from the point the traffic leaves the network (egress) to the point the traffic enters the network (ingress). This is especially true in cases where distributed attacks are launched using packets with spoofed source IP addresses.

Distributed attacks are likely to involve many source addresses, possibly from many diverse physical network paths. Near the target, traffic flows are likely to appear to be from many different source addresses and relatively few physical network paths. Near a point of origin, traffic flows may appear to be from a small number of source addresses and relatively few physical network paths. When tracing from a victim back to multiple attack sources, the traffic flows will probably deaggregate into many separate source addresses and physical network paths. The proximity of an ISP to the victim and the origin of an attack will determine the scope of an attack's traffic flow that is visible to the ISP.

Because distributed intruder systems may originate traffic from a number of different network backbones, it is likely that a global network operator will have a more complete view of the distributed nature of the attack. Smaller regional network operators are likely to see distributed attacks in aggregated form based on the number of upstream network connections.

In a distributed bandwidth denial-of-service attack, the proximity of an ISP to the end victim may have an indirect impact on the ISP and other downstream sites sharing the ISP's network resources. It is possible for portions of an ISP backbone to be overwhelmed, causing degradation and/or denial of service for sites that are not directly targeted in an attack.

Coordination among network operators and among sites involved in incidents is essential for diagnosis, tracing, and control of distributed attacks.

The following table summarizes actions the ISP community can take to better deal with distributed attacks, some actions particularly for distributed denial-of-service attacks. After the table are further explanations.

	Immediate	Short Term	Long Term
Protect	Establish crisis policy and procedures. Maintain and enforce an acceptable use policy.	Do ingress filtering. Disable directed broadcasts.	Educate customers. Implement automated anti-DoS policy enforcement.
Detect	Establish an incident response team.	Review high-profile target systems.	Automate scanning/patching of high-profile target systems. Move detection closer to the source of attack.
React	Do case-by-case egress filtering. Share information with others involved.	Establish a method for tracing back ongoing attacks to their source. Do case-by-case ingress filtering.	Establish a method for tracing back attacks in real time. Perform historical traffic flow analysis.

Table 2 – Suggestions for Internet Service Providers

Protective Measures

Immediate Actions

- Establish crisis policies and procedures.
Communicate policies and procedures to your constituency and staff. Include procedures for handling reports of attacks from the constituency and from the Internet community. Include provisions for an out-of-band emergency reporting channel in case network communication is unavailable.
- Maintain and enforce an acceptable use policy.
Include provisions to allow the ISP to track and limit service to those machines and/or networks that participate in attacks resulting from distributed-systems tools.

Short-Term (6 months) Actions

- Do ingress filtering.
Use ingress filtering to limit origination of IP packets with spoofed source addresses. The goal is to increase the ability to identify components of distributed systems.
- Disable directed broadcasts.
Prevent the use of networks in packet amplification denial-of-service attacks such as “smurf” attacks.

Long-Term (12+ months) Actions

- Educate customers.
Educate customers about potential security threats and about security best practices.
- Implement automated anti-denial-of-service policy enforcement.
Work toward an infrastructure that is able to provide automatic enforcement of policies designed to prevent denial-of-service attacks.

Detecting Attacks

Immediate Actions

- Establish an incident response team.
Pre-allocate resources to respond to security incidents.

Short-Term (6 months) Action

- Review high-profile target systems.
Establish the practice of reviewing infrastructure systems that may be highly visible targets for hosting distributed systems.

Long-Term (12+ months) Actions

- Automate the review and patching of high-profile target systems.
This automation helps to reduce the risk of having critical systems compromised due to well-known vulnerabilities for which there are patches.
- Move the initial detection point closer to the source(s) of attack.
Rather than detecting attacks close to the victim, work toward an infrastructure that makes it possible to detect attacks closer to the attack source(s).

Reacting to Attacks

Immediate Actions

- Do case-by-case egress filtering.
Apply egress filtering to identifiable packet streams to stop attacks from leaving the network backbone and to limit the immediate effects of an attack on a victim site. “Blackholing” the victim host or network might be necessary if filtering is not possible. This should usually be done only if it does not do more harm than good. It will, of course, deny service to the null-routed host or network but will probably stop the attack closer to the source and possibly restore service to other hosts or network elements.
- Share information with others involved.
Working with other involved sites and sharing information is essential to disabling an entire distributed attack network.

Short-Term (6 months) Actions

- Establish a method for tracing back ongoing attacks to their source.
Enhance your ability to trace distributed attacks back to the source(s) or ingress point(s) using existing features and tools.
- Do case-by-case ingress filtering.
Once an attack has been traced back to a source or an ingress point, use ingress filtering to prevent the attack from entering the network backbone. Filters should be tailored to stop the particular attack rather than being general anti-spoofing filters.

Long-Term (12+ months) Actions

- Establish a method to trace back attacks in real-time.
Establish a method for real-time trace back attacks traffic flows from the victim or egress point to the source(s) or ingress point(s).
- Perform historical traffic flow analysis.
Establish a method for historical traffic flow analysis to gain global visibility for identifying distributed attack systems.

Incident Response Teams (IRTs)

This section highlights issues for incident response teams to consider for detecting, responding to, and protecting against distributed attacks. Because IRTs generally collect and process incident information from a large constituency consisting of one or more large distributed networks, they play a crucial role in the detection of and response to distributed attacks.

Because of the variation among response teams, it is difficult to provide suggestions that apply to all. When developing this section, workshop participants considered incident response teams that have one or more of the following responsibilities:

1. Coordinating and distributing security information (CERT/CC)
2. Setting and implementing site security policy (serve as a corporate IRT)
3. Coordinating response to incidents (university response teams)
4. Maintaining data integrity (audit teams)
5. Protecting very large networks (large ISPs)
6. Identifying and tracking intruders (law enforcement)

Regardless of a team's responsibilities, the best protection against attacks is to be prepared. General information about incident response teams, procedures, and policies can be found in the following sources:

Handbook for Computer Security Incident Response Teams (CSIRTs), by Moira J. West-Brown, Don Stikvoort, and Klaus-Peter Kossakowski.

<http://www.sei.cmu.edu/publications/documents/98.reports/98hb001/98hb001abstract.html>

Forming an Incident Response Team, by Danny Smith

http://www.auscert.org.au/Information/Auscert_info/Papers/Forming_an_Incident_Response_Team.html

In addition, general security advice can be found on the web sites of members of the Forum of Incident Response and Security Teams (FIRST). Links can be found on the FIRST web site: <http://www.first.org/>

The suggestions below focus more specifically on attacks using distributed-systems intruder tools. The table provides highlights, and further details follow the table.

	Immediate	Short Term	Long Term
Protect	Determine chain of command. Be aware that your infrastructure may experience consequences of an attack.	Open communication channels with your constituency: provide attack signatures; encourage reporting; provide information. Encourage your constituency to implement filters.	
Detect	Develop criteria for detecting distributed-systems attacks.	Develop procedures/ algorithms for dealing with large amounts of traffic.	Develop procedures/ algorithms for handling automated incident reports.
React	Scope the extent of the attack. Escalate the priority of identifying machines acting as masters. Block traffic from known masters. Distribute information to appropriate IRTs or law enforcement.	Encourage your constituency to capture, log, & report suspicious traffic. Deploy temporary sensors such as network sniffers or intrusion detection systems.	Provide tools & methods for detecting installation of masters & daemons if possible.

Table 3 – Suggestions for Incident Response Teams

Protecting Systems

The best step a response team can take to prevent distributed-systems attacks is to raise awareness within your constituency. They need to be aware of the concept that the security of any network on the Internet depends on the security of all other networks. The widely varying implementation of security measures is what often makes a distributed attack successful.

Some of the suggestions below are not unique to distributed attacks, but as intruder tools become more distributed these issues become more important. The appropriate time frame for action depends on the mission of the IRT, so the time frames below are suggestions.

Immediate Actions

- Determine chain of command both internally for your team and externally for providers of critical infrastructure within your constituency.
This is not specific to distributed attacks but is important to understand when handling a crisis. The information should be available ahead of time to avoid delays when the IRT is working under pressure.
- Be aware that your own infrastructure may experience consequences of distributed-systems attacks, such as denial-of-service attacks, if your network or one near your network is targeted.
Consider developing contingency plans, and establish immediate, short, and long term goals to handle distributed attacks. Use the points in this section as guidelines or a starting point.

Short-Term Actions

- Open communication channels with your constituency.
 1. Provide attack signatures – Providing signatures of known distributed attacks helps members of your constituency become sensors, contributing to your successful detection, scoping, and diagnosis of these attacks.
 2. Encourage members of your constituency to report incidents – Receiving reports of attacks and anomalies is a fundamental and necessary piece of detecting distributed attacks.
 3. Distribute information about ongoing attacks – Communication about ongoing attacks needs to flow in both directions. Informing members of your constituency about significant ongoing attacks raises awareness and provides incentive for continuing to report incident data.
- Encourage constituency to implement filters (both inbound and outbound) that can stop potential attacks.
At a minimum, encourage members of your constituency to block outbound spoofed traffic, inbound traffic associated with well-known vulnerabilities that are commonly used in tools for widespread compromise and allocation of resources, and ports that are used for communication and control in distributed intruder networks.

Detecting Attacks

Immediate Actions

- Develop criteria for detecting distributed-systems attacks.

Because response teams are often in the unique position of processing incident data from one or more very large networks, they are one of the few entities capable of detecting and understanding the scope of an attack distributed across multiple networks. Thus, we encourage response teams to carefully examine data, reports of incidents, and output from intrusion detection systems looking for signs of distributed attacks. Ultimately, response teams should strive to distinguish distributed attacks from other activity.

Relying on signatures for identifying specific distributed attacks is not enough since teams receive data about new and novel tools and attacks. It is important to consider how future attacks may be detected, considering that the intruder community is moving toward distributed models for many types of tools.

Short-Term Action

- Develop procedures/algorithms for dealing with large amounts of traffic, and share them with other teams.

A problem not unique to distributed attacks is finding mechanisms to efficiently process large amounts of data received from diverse sources without missing anything important. As intruder tools continue to develop toward distributed models, it becomes increasingly important to use mechanisms for automatic processing of incident data. IRTs can benefit from sharing tools and effective algorithms for detecting distributed attacks.

Long-Term Action

- Develop procedures/algorithms for handling automated incident reports.

In the long term, a community effort is needed to develop procedures and algorithms for handling automated incident reports. An important component of that is developing a common language for representing incidents. Several efforts are under way both in the IDS community and within the CERT/CC that will enable automated incident reporting in the near future.

Responding to Attacks

Some of the distributed attacks that workshop participants have seen thus far have involved bandwidth consumption denial-of-service attacks. When responding to this specific type of distributed attack, keep in mind that resources that depend on available bandwidth (such as email) may not be reliable. In responding to attacks using distributed intruder tools, teams should take the following actions:

Immediate Actions

- Scope the extent of attack, both locally and with other response teams.

One of the most important components in determining appropriate response is finding the scope of an attack. Determining scope may require communication with multiple sites within your constituency and, often, with other response teams.

- Escalate the priority of identifying machines acting as masters.
Identifying masters is a key component of response to distributed attacks. Teams need to obtain contact information for those sites, and communicate with them to solve the problem. Depending on the situation, the optimal strategy may involve either immediately disabling masters or leaving them up to monitor and collect additional data.
- Block traffic from known masters when possible.
If it is possible, block traffic from machines known to be acting as masters. This option may be useful in situations where machines within your constituency are actively involved in an ongoing distributed attack.
- When appropriate, distribute information to appropriate response teams or law enforcement authorities.

Short-Term Actions

- Encourage members of your constituency to capture, log, and report suspicious traffic.
- Deploy temporary sensors such as network sniffers or intrusion detection systems as appropriate.

Long-Term Action

- Provide tools and methods for detecting installation of masters and daemons, if possible.

4. A Final Word

Participants in the Distributed-Systems Intruder Tools Workshop spent two-and-a-half intensive days on distributed tools and ways to address this evolving threat. This paper contains the outcome of that work. Though we have described aspects of a response for separate audiences, it is clear that coordinated action by management, system administrators, Internet service providers and network operators, and incident response teams is needed to deal effectively with the threat of these tools. To a greater extent than previously, there is a systemic cause and the need for a systemic solution as reflected in many of the recommendations in this report.

Distributed-system intruder tools demonstrate that the security of any site on the Internet depends, in part, on the security of all other sites on the Internet. Coordinated attacks across national boundaries have been observed. The tools and attacks demonstrate that a network that optimizes its technology for speed and reliability at the expense of security may experience neither speed nor reliability, as intruders abuse the network or deny its services. The intruder technology is evolving, and future tools may be more difficult to defeat.

Workshop participants encourage readers to distribute this paper widely, but also to be vigilant, keeping informed about further developments and checking web sites of organizations such as the CERT/CC, other members of the response community, and vendors.

This paper was last updated on December 10, 1999