Chapter 18

LOGICAL TRAFFIC ISOLATION USING DIFFERENTIATED SERVICES

Tinus Strauss, Martin Olivier and Derrick Kourie

Abstract This paper proposes a scheme in which the differentiated services field of IP headers is used to logically isolate network traffic for forensic purposes. The scheme is described and two example scenarios are presented to illustrate its utility. The scheme, which is based on standard networking technology, helps achieve isolation without additional network infrastructure. Moreover, the scheme is relatively easy to implement in an existing differentiated services network. The paper also discusses key design and configuration challenges that must be addressed in a successful implementation.

Keywords: Network forensics, differentiated services, traffic isolation

1. Introduction

Genge [8] describes the dilemma that faces first responders in classical (probably non-digital) events:

> "The first person on the scene is immediately confronted with a number of considerations: victims who may be in need of immediate attention, witnesses ready to melt away at the first opportunity, the possibility of further criminal activity, the responsibility of preserving whatever evidence might be remaining and securing a crime scene while maintaining safe corridors for emergency personnel. This person must weigh all these needs and make immediate decisions based on the situation. And every situation is, in some way, unique."

Current network forensic procedures often do not have the sophistication to handle real-world incidents. Frequently, a simplistic approach is prescribed: unplug the compromised network host. While this may prevent further damage to the network, it does not necessarily preserve evidence that might remain. Also, it prevents the network from supporting the operations of the organisation.

Please use the following format when citing this chapter:

Strauss, T., Olivier, M., Kourie, D., 2006 in International Federation for Information Processing, Volume 222, Advances in Digital Forensics II, eds. Olivier, M., Shenoi, S., (Boston: Springer), pp. 229–237.

Deciding whether or not to isolate a compromised host from a network [4] is a question of balancing the requirements of a forensic investigation and the necessity of maintaining the availability of resources. Disconnecting the host from the network isolates it completely, allowing the damage caused by the malicious party to be assessed, documented and possibly corrected. The state of the host at the point of its removal from the network embodies all the evidence available on the system. The evidence might not be enough to convict the perpetrator, or the criminal activity up to that point might not have been severe enough to warrant the cost and effort of prosecution.

On the other hand, if the host is left connected to the network, the perpetrator might be unaware of the fact that his activities have been discovered and continue to engage in them. These activities could then be recorded, resulting in a stronger case against the suspect. The activities may, however, cause further damage to network assets.

In a true networked environment this problem is exacerbated by the fact that an incident is likely to involve multiple hosts and the entire network may have to be unplugged. Casey [4] suggests following approach to address these concerns:

> "However, when the system is a critical component of a network, it may be necessary to involve network administrators to reconfigure a router or firewall, partially isolating the system but permitting vital connections to enable an organisation to remain in operation."

While it is possible to manually configure, or even build, custom solutions for networked systems that allow the type of forensic isolation alluded to above, a more general solution is required. The fact is that one does not necessarily know where and when an incident will happen, and a solution that can be deployed in a significant number of cases, once an attack is underway or has just occurred, is required. It is therefore necessary to achieve isolation based on technologies that are already deployed and that are well understood by network administrators.

This paper presents a scheme that provides a balance in that the host is not removed from the network, but a variable degree of isolation is achieved through the logical separation of relevant packets from the rest of the traffic. This is achieved by using the Differentiated Services scheme, a standards-based technology commonly supported in routers and other networking equipment.

The degree of isolation depends on the nature of malicious activity. If the compromised host is used as a platform for further crimes, this limited isolation with surveillance is ideal since the activities can be monitored. If the result of allowing the suspect to continue with his activities becomes too costly, the node and the suspect can be disconnected from

the network. The captured network traffic can then be analysed and presented as evidence [3, 6].

The next two sections discuss the concept of Differentiated Services and the logical isolation scheme based on Differential Services. The final section, Section 4, concludes the paper and identifies avenues for future work.

2. Differentiated Services

The Differentiated Services (Diffserv) scheme [1] was devised as a scalable approach to service differentiation in IP networks. The scheme examines the DS field [10] in the IP header and, based on the value in the field, a packet is treated in some predefined manner at each hop on the path to its destination. The value of the DS field is referred to as the differentiated services codepoint (DSCP).

Packets are marked, i.e., assigned a DSCP, according to certain rules and conditions. The marking decisions could be based on temporal properties of the arriving packet stream or on something as simple as the source of the arriving packet.

One reason for the scalability of the Diffserv scheme is that sophisticated functions such as classification and marking are only performed at the boundary of a Diffserv-capable network; the interior nodes simply use the (DSCP) marks to determine how to treat packets. Another reason is that packets are aggregated into groups or classes with the same DSCP. Traffic is considered in aggregates that require the same treatment: flows are not treated individually.

The IETF has specified two per-hop behaviour groups: the Expedited Forwarding (EF) per-hop behaviour group [5, 7] and the Assured Forwarding (AF) per-hop behaviour group [9]. EF per-hop behaviour provides a building block for low delay and low loss services; the intent is to ensure that suitably marked packets encounter short or empty queues in the forwarding path.

AF per-hop behaviour provides a means for a network operator to offer different levels of forwarding assurance to packets. Four AF classes are defined and each class is allocated resources, such as buffer space, in each of the nodes in the Diffserv network. The packets belonging to each of the four classes are thus, in a logical sense, isolated from each other. Within each class, packets are marked with one of three drop precedence values. The precedence value assigned to a packet depends on whether or not the traffic stream is within the agreed-upon profile. If the traffic stream exceeds the profile, packets are given a higher drop precedence. Packets with higher drop precedence are discarded with a higher probability than

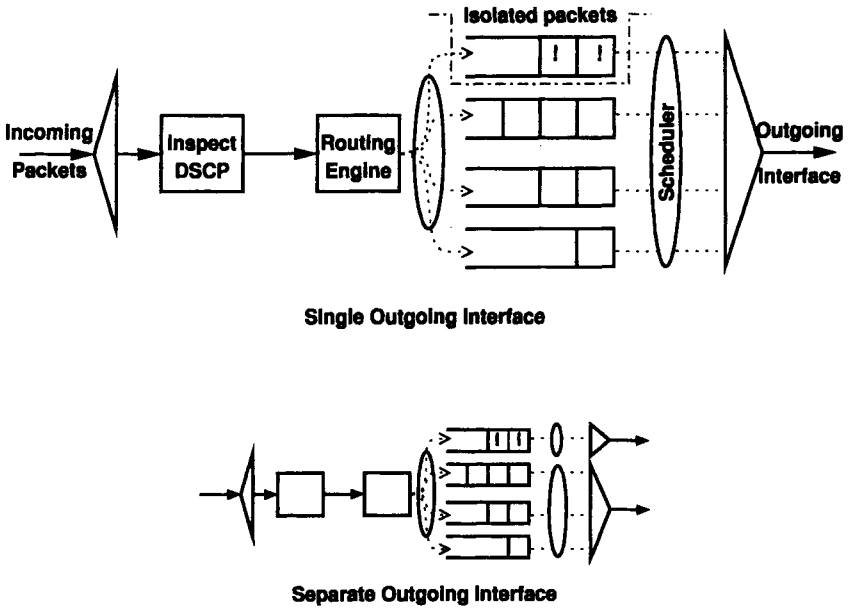

Figure 1. DSCP-based isolation inside a router.

packets with lower precedence. In this way, a network provider is able to offer different levels of service to its subscribers.

Since the provider of the Diffserv network, be it an organisation or an ISP, is free to implement its own marking policies, the organisation may reserve a set of DSCPs to assign to packets that are of forensic interest. Such a scheme is described in the next section.

3. Employing Diffserv for Logical Isolation

The idea behind the logical isolation scheme is simple: assign a DSCP to a packet that is deemed to be of forensic interest. The packet is thus identifiable throughout the Diffserv domain as associated with a specific criminal activity or surveillance effort. The marked packets are then placed into dedicated queues, which logically isolate them from regular packets.

Figure 1 provides a graphical representation of the scheme within a single Diffserv-capable router. When a marked packet arrives at an incoming interface of a router, its IP header (including the DS field) is inspected, and a decision is made about the relevant outgoing interface. The packet is then placed into the appropriate queue. If the packet is marked as forensically interesting (denoted by "!" in Figure 1), it is

placed in the dedicated queue; otherwise it is placed in one of the other queues. In this way, forensically interesting packets are logically isolated.

The scheduler can be configured to manipulate the rate at which forensic packets are serviced by the node. It is, therefore, possible to slow the traffic down.

Note that the routing decision may or may not take the DSCP into account. In Figure 1 (upper diagram), it is assumed that the routing is done independent of the DSCP. In the lower diagram of Figure 1, packets earmarked for forensic analysis are placed in a queue with a dedicated outgoing interface. This provides a greater level of physical isolation of packets.

If the DSCP is considered when making a routing decision, it is possible to route packets of forensic interest differently from regular packets. This enables the network operator to steer marked packets to certain points in the network where they may be captured for preservation and analysis. The router may be configured so that all packets matching the relevant DSCP—in addition to being forwarded to the outgoing interface—are copied and sent to a secondary interface where they are recorded.

3.1 Example Scenario

Figure 2 presents a network that implements the isolation scheme. The network is connected to the Internet, and two suspects (S1 and S2) are communicating with hosts H1 and H3, respectively. The network incorporates a marking station (MS), preservation station (PS) and management station (MGT). The marking station is responsible for marking packets; this function could be performed by an appropriately configured firewall or intrusion detection system (IDS). The preservation station collects and preserves network traffic for reconstructive traffic analysis. The management station configures and manages network elements.

Assume that host H1 is compromised and that the compromise is discovered. Instead of disconnecting H1 from the network, the investigators decide to keep it connected and monitor the situation. The Diffserv-enabled switch is configured (via the management station) to mark all packets entering the switch on the port connecting host H1. The packets are then isolated, and prevented from adversely affecting regular network traffic.

Figure 2 illustrates the case where the DSCP is also used to make a routing decision, since the packets are steered through the preservation station PS. The marking station MS at the boundary with the Internet marks all incoming packets destined to host H1 and these are steered to

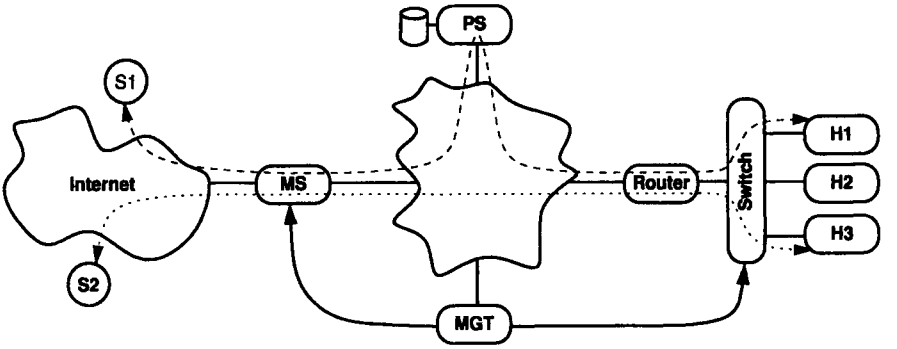

Figure 2. Example network.

the preservation station. All network traffic between S1 and H1 is now isolated and recorded for forensic analysis. Also, any traffic generated by H1 is marked and isolated, which protects the rest of the network infrastructure.

The scenario described is reactive in nature: a break-in was detected upon which a configuration change was made. It is also possible to have the IDS serve as a marking station. Traffic identified as being suspicious by the IDS is marked and automatically isolated. The traffic could then be collected or monitored in real time to determine if it is of forensic interest.

The communication between S2 and H3 in Figure 2 illustrates the case where the marked traffic is isolated but not steered to the preservation station. In this case, collection could occur through some other means.

3.2 Implementation Challenges

The isolation scheme is simple and elegant. There are, however, several challenges that must be addressed when implementing the scheme.

One challenge is deciding which packets to mark. The decision is simple if an investigation is already in progress: mark all packets headed to a certain destination or originating from some source. Alternatively, as in the example above, one might mark all the packets destined to the host as well as all the packets that enter the Diffserv-capable switch on the port that serves the compromised host. It is more difficult to mark traffic that is not yet associated with a crime, but which might be useful for forensic purposes.

The location of the marking station is also an issue. It would seem to be appropriate to place the station at a choke point in the network—at the Internet connection or at the server farm. An IDS or firewall

might be configured to mark packets that meet certain criteria instead of blocking them. Marking typically occurs at the boundary of the Diffserv domain and not at the nodes inside the domain. Multiple marking stations could be implemented at the Internet gateway, server farm or elsewhere.

Since Diffserv provides different levels of service to different packets, it is necessary to decide which service levels should be granted to the marked traffic. Two of the levels at which to consider this issue are the engineering/provisioning level and the operational level. How much capacity should the operator dedicate to the isolated traffic? Should the operator be able to adjust the isolated traffic to slow it down to reduce damage or aid in real-time analysis? It seems appropriate to minimise the loss of isolated packets in order to preserve evidence.

The location of the preservation station should be considered carefully. Since all the relevant marked traffic passes through the preservation station, the location of the station can impact the load on the network. Resource bottlenecks could result if preservation stations are placed inappropriately. Note that multiple preservation stations could be positioned at suitable locations in the network.

The scheme described here is designed for a single Diffserv domain. It is possible to extend it to multiple domains if all the domains agree on the DSCPs to use for marking traffic.

Diffserv is based on traffic aggregates, not single flows. The scheme, therefore, isolates marked traffic from regular traffic, but it does not isolate individual traffic flows under investigation.

Note that the last marking station should remove the marks on packets when they leave the Diffserv domain. Otherwise the target of the investigation might be alerted about the surveillance effort.

3.3 Advantages

The principal advantage of the logical isolation scheme is that it does not require network operators to introduce additional technology. Of course, this assumes that the operators are already using Diffserv.

The scheme is extensible and easily modifiable. Since the scheme is based on Diffserv, which is defined by the IETF, it is a standards-based approach. This allows network operators to change equipment without performing major reconfigurations to achieve logical isolation. Moreover, the scheme is based on technology that is readily available in networking equipment. It is only a matter of configuring the Diffserv domain appropriately; this results in cost savings.

Diffserv is defined for both IPv4 and IPv6 packets, enabling the scheme to be used in networks running either IP version. Finally, since only marked traffic is captured, the scheme could aid in preserving privacy while enabling investigators to obtain evidence. This assumes, of course, that only appropriate packets are marked.

4. Conclusions

The Diffserv-based scheme provides for variable levels of logical isolation while balancing the requirements of forensic investigations and the necessity of maintaining the availability of resources. Since the scheme is built upon a flexible standards-based technology that is readily available, it is relatively easy and cost-effective to implement.

The scheme assumes that only relevant packets are marked for isolation. This is, however, not as simple as it appears because packets must be marked while a crime is being committed but before the crime has been detected. This is an important issue that needs further research.

Minimising the loss of isolated packets is also an issue. Resources must be reserved to accommodate these packets: reserving too few resources results in the loss of evidence; reserving too many leads to wasted resources. Guidelines must be developed for allocating resources.

The idea underlying the logical isolation scheme can be applied to other technologies, such as MPLS [11]. Labels could be used to isolate packets by routing along label-switched paths that are dedicated to packets of forensic interest. Another technology that could naturally allow for isolation is a provider-provisioned virtual private network (VPN) [2]. Since such a VPN isolates traffic and is under the network provider's control, it can function as a VPN for forensic purposes.

References

[1] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang and W. Weiss, An architecture for differentiated services, *RFC 2475*, December 1998.

[2] R. Callon and M. Suzuki, A framework for layer 3 provider-provisioned virtual private networks, *RFC 4110*, July 2005.

[3] E. Casey, Network traffic as a source of evidence: Tool strengths, weaknesses and future needs, *Digital Investigation*, vol. 1(1), pp. 28-43, 2004.

[4] E. Casey, *Digital Evidence and Computer Crime: Forensic Science, Computers and the Internet*, Elsevier Academic Press, London, United Kingdom, 2004.

[5] A. Charny, F. Baker, B. Davie, J. Bennett, K. Benson, J. Le Boudec, A. Chiu, W. Courtney, S. Davari, V. Firoiu, C. Klamanek, K. Ramakrishnan and D. Stiliadis, Supplemental information for the new definition of the expedited forwarding per hop behavior, *RFC 3247*, March 2002.

[6] V. Corey, C. Peterman, S. Shearin, M. Greenberg and J. van Bokkelen, Network forensic analysis, *IEEE Internet Computing*, vol. 6(6), pp. 60-66, 2002.

[7] B. Davie, A. Charny, J. Bennett, K. Benson, J. Le Boudec, W. Courtney, S. Davari, V. Firoiu and D. Stiliadis, An expedited forwarding per hop behavior, *RFC 3246*, March 2002.

[8] N. Genge, *The Forensic Casebook — The Science of Crime Scene Investigation*, Ebury, London, United Kingdom, 2004.

[9] J. Heinanen, F. Baker, W. Weiss and J. Wroclawski, Assured forwarding per hop behavior group, *RFC 2597*, June 1999.

[10] K. Nichols, S. Blake, F. Baker and D. Black, Definition of the differentiated services field in the IPv4 and IPv6 headers, *RFC 2474*, December 1998.

[11] E. Rosen, A. Viswanathan and R. Callon, Multi protocol label switching architecture, *RFC 3031*, January 2001.