Towards a Denial-of-Service Resilient Design of Complex IPsec Overlays

Michael Brinkmeier and Michael Rossberg and Guenter Schaefer
Technische Universität Ilmenau
[michael.brinkmeier, michael.rossberg, guenter.schaefer][at]tu-ilmenau.de

***Abstract*—**
By monitoring the exchanged IPsec traffic an adversary can usually easily discover the layout of virtual private networks (VPNs). Of even worse extend is the disclosure if compromised IPsec gateways are considered, for example in remote environments. This revelation enables attackers to identify vital components and may allow him to compromise the availability of the overall infrastructure by launching well-targeted denial-of-service (DoS) attacks against them. In this article we present a formal model to analyze the resilience of VPN infrastructures against DoS attacks, to estimate the impact of compromised gateways, and to formalize the planning process of more resilient infrastructures.
***Index Terms*—Denial-of-Service, Availability, Virtual Private Networks, IPsec, Modeling.**

I. Introduction

For many major organizations large IPsec infrastructures [1], [2], [3] promise a more secure and yet cheaper possibility of communication than the previously used leased lines. These IPsec virtual private networks (VPNs) (Fig. 1) consist of three basic types of components: multiple *trusted "red" networks* (striped), one or more *untrustworthy "black" networks*, such as the Internet, and two or more *IPsec gateways* to connect the trusted parts securely with regards to data confidentiality, data integrity, and data authentication. The components may be connected to form complex topologies, as every single IPsec gateway may be responsible for multiple separate trusted networks and may have one or more uplinks to untrustworthy network parts. Trusted networks may be nested, for example allowing an additional protection of sensitive departments.
However, the availability and reliability of these VPNs depend on the infrastructures' resistance against denial-of-service (DoS) attacks, implying that all vital parts of the infrastructure must be protected. The identified major threats are IPsec gateways that are either situated in an untrustworthy network or that are relatively untrustworthy themselves. An example is an IPsec gateway that is placed in a foreign field office or within the network of a subcontractor. On the one hand, an adversary may observe traffic exchanged by those gateways to obtain knowledge of the identities, i.e. the outer IP addresses, of more important parts of the VPN. On the other hand, an adversary could perhaps compromise a client computer or the IPsec gateway itself, e.g., by physical actions, and start to send legitimate IPsec traffic to the vital core of the overlay, causing a DoS inside of the VPN.
A common precaution against DoS attacks is the conclusion of a service level agreement (SLA) with the responsible

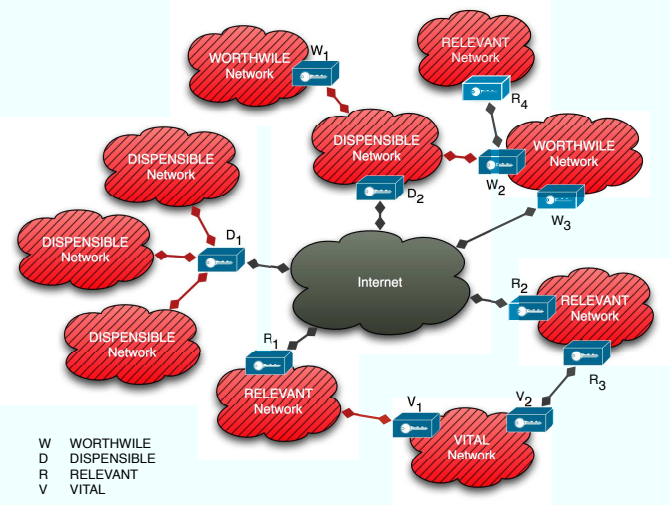

Fig. 1. A topology of the transport network in an IPsec infrastructure scenario

Internet service provider (ISP) [4], but the approach has several drawbacks: It adds bureaucratic overhead and costs. It is inflexible as the infrastructure is committed to this ISP, and a migration to other providers is not easily possible. It has an insufficient coverage as the ISP may not operate in every geographic region that parts of the VPN are located in, and even with SLAs DoS attacks may be possible, e.g., by generating forged IPsec packets that appear to come from a legitimate gateway. Therefore, no upper bound for the impact of possible DoS attacks can be calculated.
Conventional quality of service (QoS) systems [5], which give guarantees on the available network capacity, are neither suited for an impact estimation of DoS attacks as they rely on a correct router behavior. Furthermore, as they are highly dynamic, complex systems do not allow for a static formal analysis. Thus, it is very difficult to verify the availability guarantees in the presence of attackers.
For Internet environments different reactive DoS countermeasures have been proposed [6], [7], but they do not allow for a formal security analysis neither. Furthermore, some depend on the dynamic analysis of flow properties for anomaly detection. This approach is not applicable for VPN infrastructures as the flows may be encrypted and cover traffic may be used to prevent covert channels attacks.
Neither of the approaches is able to give static guarantees

for the impact of DoS attacks, which is a requirement for critical VPN infrastructures. Therefore, a static model is required to derive a guaranteed upper bound for the impact of potential DoS attacks, to rate the DoS resilience of large VPN infrastructures, and to plan infrastructures in an inherent DoS-resistant layout. Hence, in this article we contribute:

- a formal model to statically analyze and rate VPN infrastructures by giving upper bounds for the impact of DoS attacks,
- relevance classes for IPsec gateways to protect a trusted core of a VPN against DoS attacks, and
- the derivation of the requirement to route traffic between relevance classes in monotone order.

The rest of the article is outlined as follows: After introducing objectives for a model to measure the DoS resilience of VPNs in section II, a first basic model is presented in section III. This model is developed in more detail with regards to external attackers in section IV, and with regards to compromised VPN gateways in section V. Finally, we derive guidelines towards a more resilient VPN design (section VI) and give a conclusion.

II. Objectives

The need to protect the availability of the IPsec infrastructure leads to the following sub goals:

- The *resilience* of the VPN structure against DoS attacks shall be *quantified*.
- VPN components that may be *affected* by DoS attacks shall be identified.
- The *partitions*, the VPN can be split in by an adversary with certain quantifiable capabilities, shall be identified.
- VPN components that *limit the effect* of DoS attacks shall be identified.
- VPN components should be able to have different *relevance classes* to the VPN functionality (Some may be more vital than others).
- Different *attacker types* must be considered. Those are:
  - *Local observers* in untrustworthy parts of the transport network
  - *Compromised devices* in trusted networks
  - *Compromised VPN gateways* of lower relevance levels

III. A Basic Model to Estimate the Impact of DoS Attacks

A common concept for formal traffic analysis is the use of flow networks [8], [9]. The computer network is modeled as a directed graph $G(V, E)$ with every edge $\in E$ having a positive maximum capacity.

For a given VPN overlay network such a graph can easily be generated by using gateways as vertices and security associations as edges. For the sake of simplicty, we assume $V = \{g_1, \ldots, g_n\}$, i.e. the gateways are numbered from 1 to n. The capacity c_{ij} of an edge $e = (g_i, g_j)$ is the maximum throughput that can be transmitted through the underlying transport network for this association. Then the maximum flow from a gateway s to another gateway t is an upper bound of the maximum amount of traffic that the source gateway s can generate towards the target gateway t [9], [10].

The actual amount of traffic may be lower, if the transport networks capacity is reached before the capacity of the overlay network. To avoid a modeling of this transport network, we assume that the capacities of the overlay edges are set to values that do not lead to congestion in the transport network.

This model assumes that within the overlay a perfect routing is possible. I.e. a multi-path routing mechanism utilizes the maximum network capacity. The use of common routing protocols, like Open Shortest Path First (OSPF), in the VPN will lead to lower network utilization as only single paths can be used to forward packets at a time. It is important to note that this property may not hold, if a distance vector protocol is used and routing cycles exist in the network as packets may then use a link multiple times. Consequently, all analyses referring to a flow model will consider the strongest possible attacker, since we consider that the routing mechanism supports the attack.

In addition to the assumption, that the routing is perfect, the flow model makes another implicit assumption, which has to be considered carefully. In general, a DoS attacker does not mind, whether some packets are lost due to capacity restrictions. Hence, it does not limit itself to the maximum flow (or a routing that does not lead to packet losses) and tries to flood the network, instead. Due to this contradiction of reality and model, the fragmentation close to the source of the attack can be much higher than the maximum flow indicates. For example we can assume, that all links leaving the source are jammed during a real attack, but not in the maximum flow. Nonetheless, among all minimum s-t-cuts exists a unique minimal[1] one containing s. 'Behind' this cut at most the traffic predicted by the maximum flow can occur. As a consequence, the maximum flow does not give information about the actual fragmentation that occur during a DoS attack, but about 'barriers' – the minimum cuts – which ensure, that the caused traffic behind them is bounded by the maximum flow value. For more details about how the minimum cuts can be retrieved from a maximum flow, see [11].

Nonetheless, this quite simple model already allows for a basic analysis of DoS attacks, as long as the attacker is bound to the VPN connections, i.e. an attacker that compromised a computer within a trusted network, generating traffic towards other trusted networks. For example it is possible to calculate an upper bound of the traffic that a compromised gateway s can generate by flooding a specific target gateway t with messages, by computing the maximum flow from s to t.

IV. DoS Protection Measures against External Attackers

However, in order to estimate the impact of DoS attacks on the VPN infrastructure the simple model will not be sufficient

[1] 'Minimal' meaning, that it does not contain another one and that every minimum s-t-cut contains it.

if external attackers are considered. Attackers that are able to learn the external IP addresses of VPN gateways will not obey the capacities of the overlay. Instead, they may launch classical DoS attacks against one or more identified gateways from outside of the VPN, e.g. by utilizing a botnet. In the formal context, this can be interpreted as the removal of certain vertices from a network, possibly leading to a fragmentation of the VPN into multiple partitions. In graph theoretic terms, this problem is the detection of *vertex cuts*. Hence a natural measure for the resilience of a VPN against external DoS attacks is the *(local) vertex connectivity* $\kappa(s,t)$, counting the minimum number of vertices that have to be removed in order to separate s from t [12], [13]. For example, in the network in Fig. 2, we have $\kappa(R_1, D_1) = 1$, since the removal of gateway W_3 separates R_1 from D_1. Similar, we have $\kappa(D_1, W_2) = 2$.

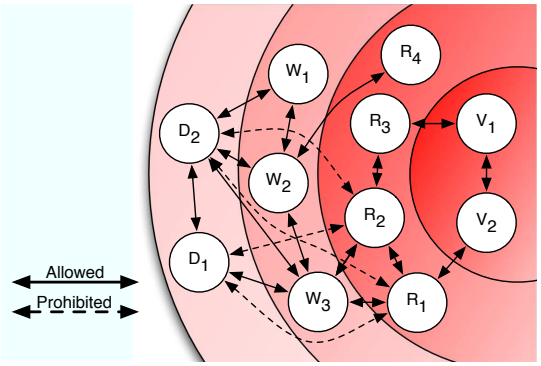

Fig. 2. The flow network for a multi-layered VPN overlay structure

In many systems, a single gateway has more than one *interface*, each one of them providing a connection to different other gateways of the VPN using another IP address. In general, a DoS attack cannot destroy the whole system, but only jam one or more interfaces. Hence, a more detailed model of the VPN requires the introduction of interface nodes. On first view, this can be done by replacing each gateway by a star, whose center vertex is connected to its interface nodes. The security associations are then edges between the two communicating interfaces. But at second thought, this model allows a direct attack on the gateway as such, implying a failure of all interfaces at once. Hence, another model seems more appropriate.

The interfaces for each gateway are modeled to form a complete subgraph, or clique, whose edges have infinite capacity. In the graph model the vertices are denoted by $g_{i,j}$, meaning the j-th interface of gateway g_i. Furthermore, we have two types of links. The *intra-gateway links* connect interfaces of the same gateway, i.e. they are of the form $(g_{i,j}, g_{i,j'})$, and they have infinite capacity. The *inter-gateway links* connect interfaces of different gateways, i.e. they are of the form $(g_{i,j}, g_{i',j'})$ with $i \neq i'$, and have the capacity restriction $c(g_{i,j}, g_{i',j'})$ of the associated security restriction. In addition, the vertex $g_{i,j}$ has the capacity of the respective physical j-th interface of the gateway g_i. An example for this construction is given for gateway W_2 in Fig. 3.

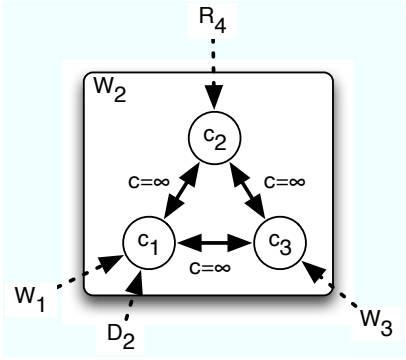

Fig. 3. Example of a flow network to estimate the impact of external DoS attackers

In order to identify the partitions a network is split into by a DoS attack, it is sufficient to remove all interfaces vertices that an attacker was able to identify. Then all partitions are calculated efficiently by finding all strongly connected components in the graph [14]. The obtained result contains all regions of the network that are still able to communicate bidirectionally in the event of the projected attack.

On the one hand, in large VPN infrastructures the approach to calculate every possible set of attacked gateways fails for reasons of complexity as an administrator may not be able to survey the influence of every possible s on every possible target t. Therefore, a reduction of the complexity is required, which can be reached by grouping gateways, and, even more important, the introduction of groups allows for a formal analysis of the possible influence of these groups among themselves.

On the other hand, large VPNs contain systems of different relevance, e.g. vital core systems like a central authentication database, and there are systems which act simply as clients, not providing services. Hence, it is more important, to protect the vital core of the system, than to ensure, that every client can connect. One way to reach this is the introduction of zones, such that zones of high relevance are 'protected' from less relevant zones.

Hence, following the concept of security labels [15], an ordered set of *relevance levels* $\mathcal{L}$ can be used. Thus, every gateway g is given a relevance level by a classification function $\mathrm{cl}(g) \in \mathcal{L}$. As an example the levels VITAL > RELEVANT > WORTHWHILE > DISPENSABLE are shown in Fig. 1. Usually, these relevance levels can be represented by a set of natural numbers, or even more general $\mathcal{L} = \mathbb{N}$.

The relevance labels should be chosen to represent the importance of the gateways for the VPN. Since in most networks, vital, or more relevant parts, are more thoroughly secured than dispensable parts, it is likely that the relevance labels correspond to the strength of intrusion prevention mechanisms implemented in the system. As a consequence, it is much more likely, that an attacker tries to compromise gateways of lower relevance level in order to harm more important systems. Therefore, we only consider DoS attacks from lower relevance

levels to higher ones.

One implication of the flow model described above, is the communication between distinct relevance labels should be limited. If, for example, a DISPENSABLE gateway can directly communicate with a VITAL one, the danger that the IP address of the VITAL node is revealed to less secure systems is quite high. Hence, we require that direct communication may only take place between gateways in adjacent relevance levels, i.e. gateway g may only communicate with g', if $|\mathrm{cl}(g) - \mathrm{cl}(g')| \leq 1$. This will cause the overlay network to form a multi-layered setup, as shown for our example network in Fig. 2.

The main consequence of this restriction is the limited influence of external observers. E.g., an observer that analyzes all traffic of a DISPENSABLE gateway will not be able to identify VITAL or RELEVANT gateways. He will not be able to launch DoS attacks against these more relevant parts of the VPN infrastructure, therefore.

Nonetheless, it may still be possible for an external attacker to fragment the VPN into multiple partitions by saturating links to other, possibly more relevant, gateways. In our example network (Fig. 1) the gateways D_2 and W_3 could be attacked by an external adversary that can observe traffic from D_1, for example. The outer interfaces of D_2 and W_3 would then be jammed, and both gateways would only be able to communicate over their internal interfaces. Therefore, the gateways $\{\mathrm{D}_2, \mathrm{W}_1, \mathrm{W}_2, \mathrm{W}_3, \mathrm{R}_4\}$ are separated from the rest. Furthermore, the gateway D_1 is separated from all other gateways as it can no longer communicate with D_2 and W_3 and a direct communication towards R_1 and R_2 is prohibited for security reasons (otherwise the attacker would be able to attack those). The third partition contains $\{\mathrm{R}_1, \mathrm{R}_2, \mathrm{R}_3, \mathrm{V}_1, \mathrm{V}_2\}$. Thus, the topology of the example VPN cannot give availability guarantees as attackers that are able to observe traffic of DISPENSABLE gateways can perform DoS attacks and separate RELEVANT gateways.

We say a VPN is resilient again external DoS attacks, if an attack on interfaces of gateways with $\mathrm{cl}(g) \leq l$ does not affect the reachability between gateways of levels $> l$, i.e. the subgraph induced by all interfaces of higher relevance is not partitioned. One way to achieve this, is to require for every $l \in \mathcal{L}$ that the subgraphs induced by all interfaces of level $\geq l$ is strongly connected. Then the removal of interfaces of lower relevance, does not affect any edge between interfaces of higher relevance.

V. DoS Protection Measures against Compromised VPN Components

The previous approach, to hide the externally routable identity of more relevant VPN components from less relevant internal, is not applicable for protecting against internal attackers as VPN gateways must be able to communicate with other gateways of all relevance levels. This leads to two possible attack scenarios:

First, compromised systems in trusted networks, e.g. client computers, are able to generate traffic towards other trusted networks. This means attackers that control such a computer may start DoS attacks against parts of the VPN, they had previously no access through the transport network to. However, the amount of traffic generated by a trusted device and sent towards another part of the VPN can be artificially limited by either filtering packets completely or perform traffic shaping in all intermediate gateways.

Second, one or more compromised gateways are able to create a virtually unlimited flow of valid IPsec traffic towards directly reachable gateways by handing the current cryptographic key to other systems that are under control of the attacker, i.e. by equipping a botnet with insider knowledge(see Fig. 4 for an example). These connections may be used to infuse more packets into the system that do seem to be originated from one of the compromised gateways. Thus, edges from an adversary must be modeled to have an infinite capacity.

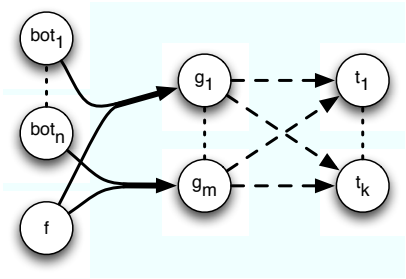

Fig. 4. DoS attack by a compromised gateway f amplifying his capacity with an external botnet

The maximal allocatable capacity per traffic flow and gateway can then be interpreted as a tensor $\mathcal{C}$ with four indices $(s, g_{i,j}, g_{i',j'}, t)$ that describe how much traffic will be allowed to flow from source s to target t if passed from gateway g_i to $g_{i'}$ using the j-th and the j'-th interface, respectively.

This approach can be modelled formally by a multicommodity flow network with commodity wise capacity restrictions. Each pair (s,t) of gateways corresponds to a commodity. These have to be routed simultaneously through the overlay network, resulting in a multiple, commodity wise flow $f^{s,t}(g_{i,j}, g_{i',j'})$ for each edge $e = (g_{i,j}, g_{i',j'})$. It has to satisfy several restrictions:

1) For every gateway pair $\{s,t\}$ and $g_{i,j} \notin \{s,t\}$:
$$\sum_{i' \neq i, j'} f^{s,t}(g_{i',j'}, g_{i,j}) = \sum_{i' \neq i, j} f^{s,t}(g_{i,j}, g_{i',j'})$$
2) For every edge $e = (g_{i,j}, g_{i',j'})$:
$$\sum_{\{s,t\}} f^{s,t}(g_{i,j}, g_{i',j'}) \leq c(g_{i,j}, g_{i',j'})$$
3) For $e = (g_{i,j}, g_{i',j'})$ with $i \neq i'$ and every gateway pair $\{s,t\}$:
$$f^{s,t}(g_{i,j}, g_{i',j'}) \leq \mathcal{C}(s, g_{i,j}, g_{i',j'}, t)$$

The two first conditions are the usual conditions for multicommodity flows, ensuring that every commodity flow is

indeed a flow and restricting the total flow over an edge. The third condition is an addition, ensuring, that the capacity restriction for each pair of gateways is satisfied. The integer multicommodity flow problem is known to be NP-complete, while the relaxed continuous version can be solved using linear programming in polynomial time [16], [17]. Even though real world applications usually would require integer solutions, the continuous relaxation provides an upper bound, leading to more conservative estimations about the DoS resilience.

In the following, we give assertions on $\mathcal{C}$, how they correspond to different properties of DoS resistance, and discuss implications on overlay network management. First, the routing protocol may be modified to better protect the VPN core. In a second step the VPN gateways may use traffic shaping to artificially limit the capacity for each flow through the overlay.

A. Relevance-aware Routing

In case the routing protocol of the overlay allows the usage of arbitrary paths, packets flowing between an attacking gateway to any other gateway may jam links of gateways that are more relevant than both communication endpoints (like the connection between D_1 and D_2 in Fig. 5). Furthermore, packets may be routed back and forth between relevance levels (the dashed connection between D_2 and V_2). This behavior is undesired, as structures of higher relevance depend on those of less relevance and such traffic may allow an adversary to attack VPN structures of higher relevance levels more easily.

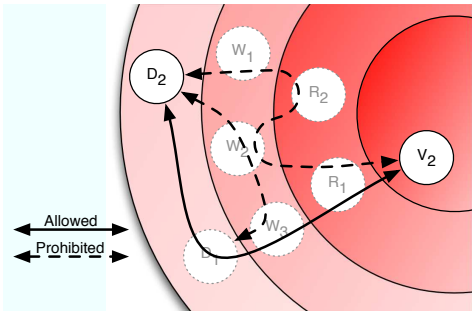

Fig. 5. Monotone flows in a relevance aware network

As a countermeasure all intermediate VPN gateways g_i enforce that the classification of next gateway $g_{i'}$ in the routing path is monotone with respect to the relevance levels, i.e.

$$\mathcal{C}(s, g_{i,j}, g_{i',j'}, t) > 0 \tfrac{bits}{s} \Rightarrow \\ \mathrm{cl}(s) \leq \mathrm{cl}(g_i) \leq \mathrm{cl}(g_{i'}) \leq \mathrm{cl}(t) \vee \\ \mathrm{cl}(s) \geq \mathrm{cl}(g_i) \geq \mathrm{cl}(g_{i'}) \geq \mathrm{cl}(t)$$

Consequently, a packet can only be routed through levels of relevance between that of the source and the target. Furthermore, a border between two neighboring relevance levels can be crossed only once, making the saturation of this border – and a resulting fragmentation of the network – hard to achieve. This rather strong requirement ensures that compromised IPsec gateways will not be able to congest any link above the level of their highest routable communication partner under any circumstances. In most scenarios an additional more fine-grained control is required.

B. Analyzing the Effects of Traffic Shaping

If the traffic is shaped by $\mathcal{C}$, as described above, the impact that an adversary s can create by sending arbitrary traffic to a target t will be bounded using the previously described multicommodity flow model. More precisely, we have to consider the directed graph $G^{s,t}$ with capacities $c^{s,t}(g_{i,j}, g_{i',j'}) = \mathcal{C}(s, g_{i,j}, g_{i',j'}, t)$. As already mentioned before, all edges from s towards directly connected gateways have to be modeled with infinite capacity since the attacker may use a botnet to infuse additional traffic into the system. A maximum flow f in this network then provides an upper bound of the maximum traffic routable from s to t.

From the residual network of $G^{s,t}$ under this flow, one can deduce information about critical parts of the infrastructure, limiting the communication from s to t. First of all, one can decide, whether the target gateway t may be separated from the remaining network, by checking wether all of its incoming links are saturated. Secondly, all minimum s-t-cuts can be reconstructed from the residual network of $G^{s,t}$ under f, leading to information about the possible fragmentation of the network [11]. This allows to identify security associations and interfaces, whose capacity is completely used by *every* maximum flow from s to t. For example, a saturated interface vertex shows that the physical constraints of a gateway limit the impact of a DoS attack, while a saturated security association is caused by the traffic shaping. The identified gateways can then be further protected or the capacities of the identified edges can be adapted for a further limitation or relaxation of the allowed flow.

For more advanced attack scenarios, in which many compromised gateways $s_1, \ldots, s_l$ attack many targets $t_1, \ldots, t_k$, the situation becomes more complex as we have to consider more than one commodity, namely those of the type (s_i, t_j).

But since we only use upper bounds for the possible traffic, we may even simplify further, leading to a standard flow problem again. Instead of assigning a capacity for each commodity to each edge (and one global capacity), we may simply think of them as one commodity, for which the flow on the edge $(g_{i,j}, g_{i',j'})$ is bounded by its physical capacity $c(g_{i,j}, g_{i',j'})$ and the total sum of all security associations, leading to

$$\min \left\{ c(g_{i,j}, g_{i',j'}), \sum_{(s_x, t_y)} \mathcal{C}(s_x, g_{i,j}, g_{i',j'}, t_y) \right\}.$$

In addition an artifical source s and an artificial sink t have to be added, such that s is connected to each s_x by an infinity capacity edge, while every t_y is connected to t by such an edge. Then a maximum s-t-flow is an upper bound for the traffic that can be generated by the compromised gateways s_x towards the gateways t_y.

VI. Guidelines

From the model and the above considerations, we can deduce some guidelines, how a DoS resilient VPN should be designed.

- One should introduce a totally ordered sequence of relevance levels $\mathcal{L}$ and ensure, that gateways communicate only with gateways in adjacent levels. This feature ensures that the observation by an external attacker only reveals gateways of the next higher level, limiting its influence on the more relevant core.
- The traffic inside the VPN should be routed from s to t so that monotone in- or decreasing relevance levels are ensured and that it can be guaranteed that the communication between high relevance gateways does not depend on less secure gateways. Furthermore, the monotony of the routing ensures that a DoS attack only passes borders between levels in one direction, reducing the caused separation of the levels.
- The higher relevance levels should be designed to be strongly connected, independent of the lower levels. This ensures that the core network still provides service, if attacked from outer levels.
- The traffic shaping should be deployed, ensuring that bottlenecks, i.e. minimum cuts to higher levels, are moved as close as possible to the outer levels. This might be realized by making the level borders to minimum cuts between vertices of lower relevance and those of higher relevance, guaranteeing that a DoS attack from lower levels is blocked to a certain degree by the border level. Additionally, the capacity crossing higher borders should be greater than that crossing lower borders. In this way, it can be guaranteed, gateways of higher levels can communicate with each other even during a DoS attack.
- Generally, gateways of the same level should be as highly connected as possible. This includes vertex- as well as edge-connectivity.

VII. Conclusion

Within this article a novel approach for the modeling of complex IPsec VPNs has been presented, allowing for a formal analysis of the infrastructures with regards to availability threads. The introduced relevance levels support administrators in planning and operating large networks, automating possibly error-prone manual tasks. The derived guidelines are simple to follow rules that restrict the effect of DoS attacks effectively. In the future, we plan to use the proposed model as a basis for further DoS impact estimations, and research on suitable approximations schemes for the particularly occurring multi-commodity flows. Furthermore, a graphical editor can ease the understanding and the planning of capacity tensors for IPsec infrastructures and automatically estimate the impact of decisions. From a theoretical point of view, the extension towards multicast scenarios seems an interesting aspect, requiring to extend the graph model to a hypergraph model. Finally, it is possible to reverse the relevance model to estimate how much data may flow out of trusted part of the network into a less trusted one, and therefore quantify the data leaks that are created by a compromise.

References

[1] V. Bollapragada, M. Khalid, and S. Wainner, *IPSec VPN Design*. Cisco Press, 2005.

[2] Cisco Systems, Inc., "Dynamic Multipoint VPN (DMVPN)," 2006.

[3] Cisco Systems, Inc., "Cisco Group Encrypted Transport VPN," 2007.

[4] R. Ramanujan, M. Kaddoura, J. Wu, C. Sanders, and K. Millikin, "VPNshield: protecting VPN services from denial-of-service (DoS) attacks," in *DARPA Information Survivability Conference and Exposition, 2003. Proceedings*, vol. 2, 2003, pp. 138–139.

[5] M. Menth, S. Kopf, J. Charzinski, and K. Schrodi, "Resilient Network Admission Control," *Computer Networks*, 2008.

[6] Y. Xu and R. Guerin, "A Double Horizon Defense Design for Robust Regulation of Malicious Traffic," in *Securecomm*, 2006, pp. 1–11.

[7] M. Waldvogel and T. Köck, "Light-weight End-to-End QoS as DoS Prevention," in *32nd IEEE Conference on Local Computer Networks, 2007. LCN 2007.*, 2007, pp. 246–248.

[8] J. Kong, M. Mirza, J. Shu, C. Yoedhana, M. Gerla, and S. Lu, "Random flow network modeling and simulations for DDoS attack mitigation," in *IEEE International Conference on Communications*, 2003, pp. 487–491.

[9] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows*. Prentice Hall, 1993.

[10] L. R. Ford and D. R. Fulkerson, "Maximal flow through a network," *Canadian Journal of Mathematics*, vol. 8, pp. 399–404, 1956.

[11] J.-C. Picard and M. Queyranne, "On The Structure of All Minimum Cuts in a Network and Applications," *Mathematical Programming Strudy*, vol. 13, pp. 8–16, 1980.

[12] R. Diestel, *Graphentheorie (Graph Theory)*. Berlin, Heidelberg, New York: Springer, 2000, 2. Auflage.

[13] M. R. Henzinger, S. Rao, and H. N. Gabow, "Computing Vertex Connectivity: New Bounds from Old Techniques," in *FOCS*, 1996, pp. 462–471.

[14] R. E. Tarjan, "Depth-First Search and Linear Graph Algorithms." *SIAM J. Comput.*, vol. 1, no. 2, pp. 146–160, 1972.

[15] E. G. Amoroso, *Fundamentals of Computer Security Technology*. Prentice-Hall, Inc., 1994.

[16] S. Even, A. Itai, and A. Shamir, "On the Complexity of Timetable and Multicommodity Flow Problems," *SIAM J. Comput.*, vol. 5, no. 4, pp. 691–703, 1976.

[17] L. Fleischer, "Approximating Fractional Multicommodity Flow Independent of the Number of Commodities," *SIAM J. Discrete Math.*, vol. 13, no. 4, pp. 505–520, 2000.