Automatic Creation of VPN Backup Paths for Improved Resilience against BGP-Attackers

Michael Grey Michael Rossberg Guenter Schaefer

Ilmenau University of Technology

firstname.lastname@tu-ilmenau.de

ABSTRACT

Virtual private networks (VPNs) play an integral role in corporate and governmental communication systems nowadays. As such they are by definition an exposed target for attacks on the availability of whole communication infrastructures. A comparably effective way to disturb VPNs is the announcement of the involved IP address ranges by compromised BGP routers. Since in the foreseeable future criminals may focus on such attacks, this article discusses the intelligent creation of backup paths in the context of VPNs as a countermeasure. The proposed system is evaluated in simulations as well as in a prototypic environment.

Categories and Subject Descriptors

C.4 [**Reliability, availability, and serviceability**]; C.2.3 [**Network management**]

Keywords

Denial-of-Service, Availability, Backup Paths, BGP, Virtual Private Networks

1. INTRODUCTION

With the emergence of the Internet, efficient global communication reached nearly every spot of our daily lives and became a major criterion for the success of companies and governmental agencies. In order to also use the public networks for the exchange of private and possible sensitive data, organizations usually depend on virtual private networks (VPNs), which protect all transferred datagrams between communicating entities by cryptographic tunnels. However, in contrast to confidentiality, integrity, and authenticity of the transmitted data, the availability is typically not regarded by the VPN implementation itself. Instead, service level agreements (SLAs) with the involved Internet service providers (ISPs) shall assure that attackers cannot simply perform denial-of-service (DoS) attacks. Nonetheless, even with SLAs the protection against DoS attacks is difficult: According to Forrester Consulting [9] about one third of the interviewed IT professionals state that regardless of the DoS protection strategies, outages could not be prevented.

SPONSORED BY THE

Federal Ministry
of Education
and Research

SAC'12 March 25-29, 2012, Riva del Garda, Italy.

A possible mitigation path is the usage of dynamically reconfiguring VPN approaches [18] that are able to repair associations in the case of DoS attacks, i.e., by allowing administrators to quickly relocate affected systems to different IP address ranges or redirecting traffic over different parts of the network. Furthermore, the unsuccessful attacks of the hacker group “Anonymous” on the Amazon compute centers [13] showed that conventional DoS attacks can be coped with, if precautions are taken.

Nonetheless, against very sophisticated DoS attacks, as happened for example during the Russo-Georgian War [8], even well protected systems can be expected to fail. We project that in future conflicts, governments and well-trained individuals will attack the inter-domain routing in order to perform such large-scale DoS attacks, as this currently allows huge damage with minimal effort. Like depicted in Fig. 1 the attackers may maliciously attract traffic in some regions of the network by compromised border gateway protocol (BGP) speakers.

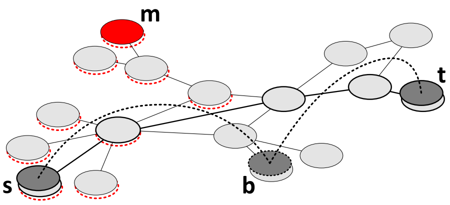

Figure 1: Attacker m attracts traffic towards t, so that s and t may only communicate indirectly via b (dotted nodes are affected by m).

In order to provide a protection mechanism for VPN against BGP attacks, we make the following contributions within this article:

- we categorize BGP attacks by their effect on VPNs,
- derive a backup path routing with negligible overhead to allow for a quick recovery from BGP attacks,
- and evaluate the mechanism by an implementation in a self-configuring VPN [18] system.

The remainder of this article is structured as follows: We will give a short introduction on BGP attacks and considerations for a backup path mechanism in Section 2. The related work in Section 3 is then compared with the given objectives. Section 4 contains details on our approach, which is subsequently evaluated in Section 5. The article concludes in section 6.

2. ATTACKER MODEL & SYSTEM OBJECTIVES

Attackers that are external to a VPN may pursue one or more of the following goals:

- **Sinkholing** refers to the sending of false route advertisements/enforcements that aim at attracting traffic to a particular site. In the context of VPNs, this allows to perform traffic analyses and to prepare grey- or blackholing attacks, where either some or all packets of a victim are dropped. As this article focuses on availability issues, we will concentrate on grey- and blackhole attackers.
- **Redirection attacks** force traffic flows to take different paths. The new intermediate nodes may also be under the control of the attacker or the enforcement might lead to a congestion collapse on specific links. Thus, this mechanism will also lead to grey- and blackholing effects.
- **Instability:** Frequent advertisements and withdrawals for networks may cause outages, as they trigger route dampening at surrounding BGP speakers. This may lead also to long convergence periods [11] as the overall BGP traffic volume increases. Furthermore, at least within BGP, attackers might try to trigger route changes that will lead to permanent loops, which need manual recovery [14]. Very VPN specific is another problem: Due to the involved asymmetric cryptography, the handshaking for a single security association may take up to several seconds, at least if smartcards are used. Thus, instabilities caused by BGP attackers may have even stronger side effects for VPNs.

Within this article, we assume an attacker that controls one or more BGP speakers within his vicinity, i.e., a close topologic distance. To create black-, greyholes, or instabilities the following strategies are possible:

- **Malicious Route Attraction:** As BGP updates are currently not protected in any way, attackers may announce maliciously short routes to a victim autonomous system (AS) or impersonate the victim AS completely. Both strategies will convince ASs that are closer to the attacker than the victim AS to route via the attacker.
- **Malicious Route Flapping:** By continuously sending malicious updates and withdrawals, further instabilities may be induced, e.g., by the creation of loops or instable routing states that BGP cannot recover from [14]. Just like the route attraction attack, attacks on the stability have commonly only local impact.
- **Prefix Hijacking:** Without further protection strategies, BGP allows attackers not only to announce fake AS information, but also arbitrary IP address prefixes that seem to be available by the AS of the attacker's choice. In contrast to the announcement of AS information, this method is not only more selective, but allows to deceive routers Internet-wide, if a more specific IP prefix is used. In order to attack a larger IP prefix, an attacker may deaggregate it into several smaller address blocks, and thus trick other ASs that a better matching IP address block is available.
- **Congestion Induction:** Attackers might use malicious routing updates and withdrawals to create congested links. While this would usually only lead to a degraded performance, a heavy congestion on links may cause a breakdown of TCP-connections. This again is known to also affect BGP peering sessions [7], causing BGP traffic spikes on session recovery and possibly avalanche effects.

In contrast to other applications, the situation seems to be more relaxed with the use of VPN infrastructures as for example impersonation and eavesdropping are prevented. Furthermore, the influence of prefix hijackings can be limited, if a /24 prefix is announced for the VPN as most BGP routers solely accept subnetworks of this minimum size.

However, as long as BGP lacks appropriate security mechanisms, the most intuitive attack mechanism in form of the announcement of a false path towards a specific destination network cannot be coped with at ease: Even VPN infrastructures suffer from this type of attack, and black-, greyholes as well as instabilities may be a consequence.

In order to cope with BGP problems, the article discusses the creation of backup paths via other VPN participants. To do so, several objectives have to be met for global VPNs:

- **No BGP details from ISP:** While for some sites it may be feasible to have access to the BGP information of the ISP, i.e., by looking glass servers, for global VPNs this can be considered impossible for all sites. Thus, a mechanism should not rely on this information.
- **Scalability:** VPN infrastructures can reach sizes of thousands of endpoints; hence, suitable mechanisms must scale over the number of participants [10] and may use local knowledge only.
- **Robustness:** The creation of backup paths shall not only increase the robustness against attackers, but also the robustness to overlay partitioning as well as general link and routing failures.
- **Security:** To strengthen the overall availability of the overlay network, backup paths shall increase the resilience against BGP attackers. However, this functionality must not create new ways to attack the VPNs.

The following will discuss existing approaches for reliable overlays and to what extend they fulfill the requirements.

3. RELATED WORK

The idea of performing an overlay routing to improve connection characteristics was first published within the *Detour* project [19]. In [20] the authors draw conclusions on various end-to-end effects of BGP-based routing in wide area networks, primarily based on own measurements and analyses of Paxson's [16, 17] prior studies. One essential part of the study is a performance comparison between both direct and alternate indirect paths. It was figured out, that for most of the direct paths, there is at least one indirect path, which is superior in sense of latency, packet loss, or available bandwidth. As another part of the project, a framework [5] for indirect routing using IP-in-IP was presented.

Somewhat similar to Detour in objectives and results are *Resilient Overlay Networks* (RONs) [2]. Among other things, the proposed algorithm tries to increase robustness and performance by making use of indirect overlay paths. This is achieved by maintaining virtual paths to all other overlay nodes, across which the connection characteristics are determined by periodic active probes, measuring for example packet loss and latency. The collected results are shared with other overlay nodes and are used for routing decisions. Within a subsequent work [4], the authors state that the incorporation of indirect paths increases the robustness of wide area network communication. However, indirect routing based on collectable properties is only convenient as long as failures on different paths do not correlate – but this frequently-used assumption seems to be untenable in typical wide area scenarios [3]. Furthermore, due to the maintenance of a full-mesh, the size of RONs is limited to about 50 nodes [21].

Moreover, several concepts were presented during the last years that focus the disengagement of transport network estimations from the actual overlay. To take only one example, the authors of [15] propose an additional abstraction layer – the so-called *routing underlay* –, which may be probed by overlay applications and allows for application specific routing decisions. The fairly scalable sharing of local inter-AS topology views serves as one primary application example. However, even if information on inter-AS topologies exists at the BGP routing tables of ISPs, they are usually not freely available to safeguard the ISP's economic interests. Even if the concepts of routing underlays promise robustness gains, the required support from (multi-homed) ISPs is not tenable.

4. SELECTING GOOD BACKUP PATHS

Bringing it all together, the exposed BGP issues – regardless of whether caused by malicious activities or not – can affect VPNs negatively. To address this problem, we present a backup path mechanism, which establishes few additional VPN associations on the basis of AS path structures to proactively protect against path outages.

4.1 Scalable Construction of Backup Paths

In contrast to approaches like RON, the creation of backup paths for global VPNs must not only focus on few additional connections, but in order to maintain scalability, another pervasive restriction must be regarded for structured VPN overlays: Only local knowledge on already established overlay connections is available and only for these the corresponding transport network structures can be determined, as the exchange of topology information must be avoided for the sake of efficiency, scalability, and possibly the resistance against byzantine attacks. This paradigm contradicts to classic backup approaches, which sacrifice the latter properties, e.g., RON and Detour periodically probe all nodes, significantly restricting the maximum overlay size.

In order to still find good backup paths, we rely on the discovery of divergent paths within the transport network, i.e., paths that are routed through different ASs. This metric can be measured fairly simple in practice, as the required information can be determined by periodically probing with *traceroute*-like tools. The obtained IP addresses of the traversed routers must then be mapped to a corresponding AS, but in contrast to approaches that rely on BGP looking glasses, these lookups do not require dedicated resources of a particular ISP. Instead online or offline databases may be queried by any VPN end-system, as this information does not change frequently. Within our approach, these active traces are exclusively performed along already established direct overlay links for scalability and efficiency reasons.

Besides relying on path probes, several other initial design decisions were taken:

- **Backup Application:** While it might be conceivable to protect only highly relevant connections, e.g., heavily utilized or structure-enforcing connections, such a differentiation is application specific and not in the scope of this article. It is assumed that one backup path is created for every direct overlay connection.
- **Candidate Pool:** For a specific overlay connection from a source s to a target t, every directly connected overlay neighbor of s represents a potential backup path *mediator*, which may be instructed to relay data of s towards t in the case of a failure. In structured overlay networks, just like the VPN auto-configuration system in [18], this typically means that every node has about $\mathcal{O}(\log n)$ possible mediators to choose from, where n denotes the size of the overlay.
- **Single Indirection:** Backup paths to circumvent local network failures might be built over multiple overlay nodes. However, within this article we will restrict on single indirect backup paths, i.e., paths with one additional hop between source and destination, for the following reasons: First, due to the restriction to solely rely on local knowledge, the creation of a multi-hop backup path will lead to difficult to assess situations and thus probabilistic decisions, as it is only possible to probe the transport network path towards a single next overlay hop. Second, a single indirection stage is to be preferred for practical reasons, because it significantly eases the stabilization and maintenance of backup paths within dynamic overlay systems. Third, the effectiveness of multiple indirections can be considered to be comparably low as can be inferred from the discussion in [2].

Under the given circumstances, it is conceivable to only rate the candidate paths based on the first part towards the backup path mediator, and leave the connection from the mediator to the reference path's destination unrated. Thus, the essential strategy of the proposed algorithm is a search for AS-divergent path segments at the first stage. However, within typical wide area network topologies this simple strategy will generally lead to good overall backup paths as will be shown in the evaluation section.

4.2 Rating Path Divergence

One particular problem is the definition of an adequate distance metric for the "best" path selection. To compare two paths, we take the classic edit-distances as a basis and calculate the longest common subsequence (LCS) [6] between the two paths; whereas the input alphabet simply consists of AS-identifiers and the AS sequences are interpreted as strings. This leads to the following recursive formula for a distance $d_{LCS}(|X|, |Y|)$ between two path representations X and Y:

$$d_{LCS}(i,j) = \begin{cases} 0, & \text{if } i = 0 \text{ or } j = 0 \\ d_{LCS}(i-1, j-1) + 1, & \\ & \text{if } i, j > 0 \text{ and } x_i = y_j \\ max(d_{LCS}(i, j-1), d_{LCS}(i-1, j)), & \\ & \text{if } i, j > 0 \text{ and } x_i \neq y_j \end{cases}$$

The LCS is merely a complementary weighted version of the well-known Levenshtein distance [12], where equality is given the weight 1 and necessary replacements, insertions as well as deletions are weighted with 0 respectively. In consequence, a backup path with small overall LCS weight is desirable, as it is an indication that different AS are passed. In contrast to more naive rating approaches, e.g., just counting the number of equal AS identifiers, the LCS is particularly suitable for the following reasons:

- LCS allows to rate not only straightforward equality, but also sequence similarity even if interrupted at one or more spots. In reality, this may occur due to specific BGP policies for the destination AS.
- As depicted in Fig 2, the distance increases once for each transposition of an AS sequence (in reference to a compared sequence). A different AS order, indicates different communication links and should be rated better than paths with an in-order AS sequence and worse than paths with completely different AS sequences.
- Attackers that artificially increase the path lengths to specific candidate mediators cannot raise their LCS rating, even though the proportion of different AS identifiers might increase.

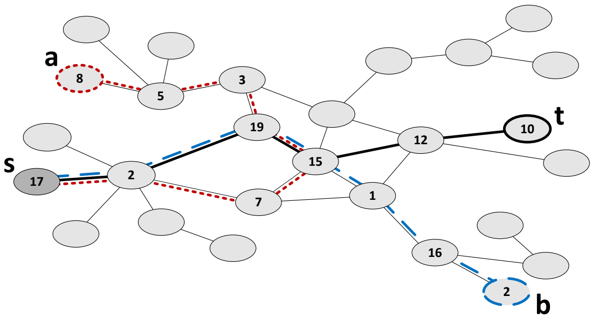

Figure 2: Transposition of AS nodes on compared paths: Given a reference path between s and t, the path between s and a is better-rated than the path towards b. Even though all paths share a dedicated subset of nodes $(17, 2, 19, 15)$ the path towards a is slightly better-rated due to the reverse traversal of $19 \rightarrow 15$.

Despite being an important indication, the sole use of LCS is not sufficient. Referencing path lengths, we can differentiate between the following situations:

- The length of a potential backup path is greater or equal to the reference path length.
- The potential backup path is shorter than the reference.

While the former case does not exert bad influence on divergence rating, the calculation suffers from the latter case as the inaccuracy rises with shorter candidate path lengths. This becomes highly relevant for very short paths to potential mediators, because in typical inter-domain topologies short paths indicate bad candidates for two reasons: First, typical wide-area networks exhibit a dense core region as well as sparsely populated regions at the edge, which usually consist of tree-like structures at autonomous system layer. Therefore, short paths between two nodes indicate similar uplinks to the core network, even if both nodes are part of different ASs, tentatively restricting the path's resilience against attacks. Second, close-by ASs are more likely to be affected by the same BGP attack. Thus, the effectiveness of close-by mediators can be expected to be exceptionally bad.

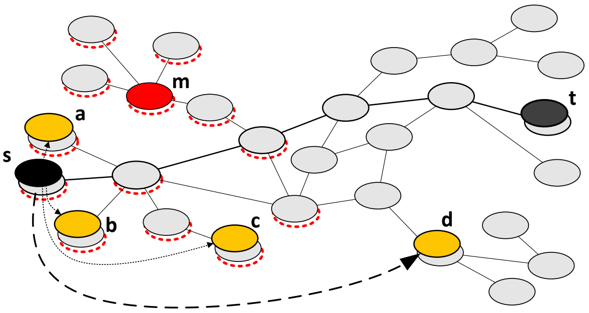

Figure 3: Candidate path length influence on backup path rating quality.

Fig. 3 illustrates the influence of very short paths: Backup paths for a reference connection between the nodes s and t may be constructed over one of the nodes a, b, c, or d, respectively. Assuming shortest paths in the transport network, the reference path shares one transport network node with each path to one of the mediator candidates. While a, b, and c are located at the same edge region as s, the path to d traverses the dense core network and therefore node d represents a considerably better mediator. An exemplary malicious AS m, that announces to be the AS of t, as well as the affected ASs are highlighted. Already this simple example demonstrates why mediators with similar topological locality should be avoided.

Unfortunately, a suitable minimal AS length depends highly on the topology of the transport network – especially on the average path length and distance to the core network – and an appropriate threshold length cannot be generically determined. Despite that, when considering a designated network type (such as the Internet), we can give an approximate value L_{Thresh}.

4.3 Backup path modes

In order to comply with different scenarios, allowing for an appropriate weighting of additional robustness and applied overhead, we distinguish between the following modes:

- **Asymmetric mode:** Independently of each other, the source and destination of a direct overlay connection choose individual backup paths. Generally this leads to two different backup paths for each dedicated link.
- **Symmetric mode:** Only one of a reference path's endpoints forces the construction of a backup path, while the other one anticipates to the chosen path. In order to select a unique coordinator, both peers use a hash-function H to calculate $V_1 = H(ID_{own}||ID_{otherSide})$ and $V_2 = H(ID_{otherSide}||ID_{own})$. Both values are compared locally and the peer with $V_1 > V_2$ has to construct the backup path. This ensures that every peer is burdened equally.

In either mode the construction of backup paths is invoked on local knowledge.

4.4 Overcoming pitfalls

After an appropriate choice has been made, the new mediator has to be informed in order to maintain a fallback-connection to t. In practice, we use a simple request-response-protocol to coordinate with the mediator, but there are some pitfalls to cope with:

First, overlay connections that are solely established by the mediator to serve as a backup path, must not be protected by the backup path mechanism. Otherwise, a rather large number of connections may be created, possibly leading to a full meshed network.

Second, it is necessary to give chosen mediators a chance to reject the construction of a backup path, especially in order to avoid overlay-loops. Fig. 4 illustrates a problematic situation, where the nodes a and b try to install backup paths over each other in order to protect their direct paths towards t. If the direct transport network connection fails, e.g., due to the appearance of an attacker m, node a recognizes an outage of the connection towards t – as node b does. Both nodes will then simultaneously fall back to their respective backup path, causing a routing loop between a and b.

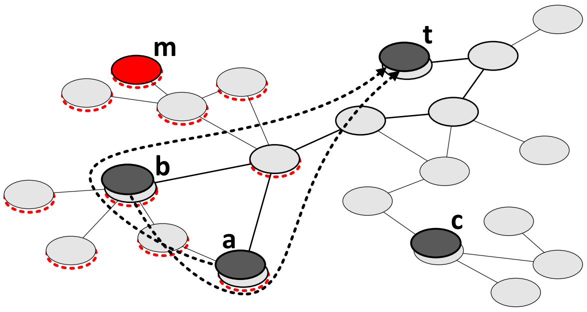

Figure 4: Potential loop caused by backup path constructions

Third, if a mediator uses a connection for itself and as a fall-back connection for another node, it may have already a different backup path for it or create one later on. This may happen in scenarios like Fig. 4, when either a or b chose c as a mediator to t. In this case, a failure may lead to a chain reaction, whereas it may happen that a valid path is established, which is redirected via multiple other mediators first. While not posing a problem, such events should be considered in an implementation. Note, that – despite the previously described simple loop protection – allowing multiple indirections may result in loops in some rare cases, however, the chances are negligible so that multiple indirections can be considered valuable without further loop protection strategies.

Forth, established backup paths need to be refreshed periodically in order to cope with the dynamics of structured overlays, ensuring the currentness of routing entries and thus the quality of the backup paths.

5. EVALUATION

To assess the proposed mechanism, we will first discuss the fulfillment of the objectives from Section 2. In a second part the increased robustness is judged by a simulative study.

- **No BGP details from ISP:** In order to rate backup paths, no access to a specific ISPs looking glass server is required. Instead the IP addresses that are obtained from the traces must be mapped to AS numbers. However, such a lookup database can be cached from a central position and is comparably static, even if failures are present.
- **Scalability:** At most one backup path of constant length is established for every overlay link and as only local knowledge is used to rate different candidates, the overhead will increase at most constantly with a growing network size. Thus, the system will scale over the number of participants.
- **Security:** In the sketched VPN scenario, the security of the overall system is an essential asset. As messages are only exchanged with neighboring nodes through cryptographic tunnels, the objective boils down to an unexploitability of the backup paths.

  Due to the proposed LCS metric, the selected mediators of our backup path mechanism cannot easily be predetermined by an external attacker, as the selected ASs are located within different parts of the Internet. This would not have been the case, if mediators close to the destination were chosen, for example.

  Furthermore, when considering internal attackers, nodes may try to fake answers during traceroutes in order to be chosen as mediators more often, and thus attract more traffic in the event of failures. However, if probes include random numbers (nonces), potential attackers can only extend the traces by inserting faked routers. But due to the used LCS metric, such an insertion cannot result in a higher rating, and hence will not lead to an exploitable situation.

Within the following quantitative evaluation, some selected questions, which were posed during the article, are examined with the help of simulative experiments:

- **Connectivity:** It will be determined to what extend the proposed backup path system enables immediate communication recovery in presence of BGP attacks.
- **Overhead:** The major cost factor in reference to scalability and efficiency is represented by the additional, proactively established VPN connections. Thus, the influence of the proposed system on the node degree is shown.

All of the experiments were performed utilizing the CAIDA [1] AS-level graph, which is based on a snapshot of collected BGP information. For the sake of simplicity, shortest paths between participating transport network nodes are statically determined for every simulation run. In order to simulate a worldwide operating organization, the VPN nodes were randomly placed at the edge of the transport network topology, forming groups of exponentially distributed size. The exponential distribution reflectse typical scenarios that contain large company sites at few ASs as well as several spatially distributed smaller sites and road warriors. Finally, to simulate malicious activities, a BGP attacker is periodically placed at a random edge node of the transport network, targeting a set of VPN nodes and the corresponding AS nodes, respectively. As motivated in Section 2, we simulate fake path announcements towards the latter ones: Thus, ASs whose path towards the attacked network is longer than the path towards the attacker's network continuously distract traffic that was originally designated for the attacked network. Due to aiming at a set of VPN participants the used attacker model is quite strong, which is assumed to attack the AS of five VPN nodes in all cases. Thus, in most cases not only one AS is targeted, but rather a set of ASs. All experiments were performed 32 times and 95% confidence intervals are shown in each of the following figures.

5.1 Impact on the Connectivity

The most prominent effect of the proposed concept is the prevention of connectivity loss, which cannot be avoided by reactive mechanisms: In case of a successful immediate backup path fallback, the overlay connection does not suffer from connectivity loss. First of all, the impact of the previously mentioned restriction on single-indirect backup paths in dependency of VPN size is examined.

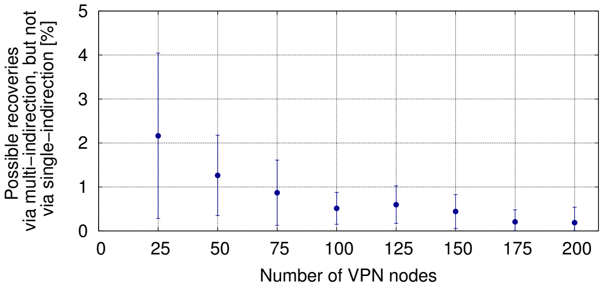

Figure 5: Situations without possible single-indirect backup path solution, but multi-indirect one

Fig. 5 shows the mean values of runs, where a possible overlay paths exists that could have prevent connectivity loss, but no single-indirect path would have sufficed. Clearly, the portion of these situations becomes negligible as the VPN size increases and it affects less than 5% of the connections even in very small scenarios. Thus, a possible benefit of multi-indirection involvement can be considered low in any case.

To evaluate the connectivity loss prevention of the proposed backup modes, an adequate L_{Thresh} must be determined (as mentioned in section 4.2). Luckily, as depicted in Fig. 6, the influence of L_{Thresh} is considerably low in CAIDA topologies. The major reason for this is the high path diversity of the Internet AS topology, which does not exhibit tree-like structures at the edges (in contrast to router topologies for example). Thus, the L_{Thresh} can be set to a value of 3, which should also make it difficult for attackers to attract the traffic of multiple close by VPNs to circumvent the backup path mechanism.

To determine the effectiveness of the approach in comparison to a random selection of mediators, another experiment was conducted. Fig. 7 illustrates the successfulness of the asymmetric mode in contrast to the randomized approach and the maximum amount recoverable by single indirection. As already discussed the latter one grows with the number of VPN nodes as more theoretic opportunities to avoid attacked areas with single indirections are created. However even if more theoretic mediators exist, the performance of the approach fluctuates at around 60% in all situations. This can be attributed to the strong attacker, which with increasing VPN size aims at more ASs on average. Furthermore, the number of nodes to choose appropriate mediators

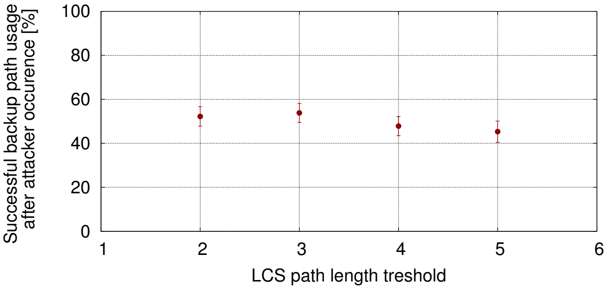

Figure 6: Situations without possible single-indirect backup path solution, but multi-indirect one.

from increases only logarithmically, due to the used overlay topology. Thus, the chance to find better mediators does not increase significantly. Interestingly, the difference between the randomized approach and the backup path mode decreases with the VPN size. This is on the one hand related to the fact, that the randomized approach suffers from choosing mediators close to the destination or source AS, which is especially problematic in VPNs of small size. On the other hand, due to the already stated high path diversity in the Internet topology, the randomized approach comes closer to the backup mode, but never performs better. In any case, the difference can be awaited to be much more significant with an increasing correlation of transport network paths, i.e., if networks are small and few ISPs are used, which makes the LCS-based choice more suitable for many VPN scenarios.

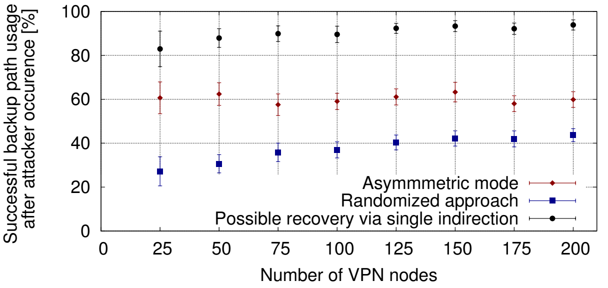

Figure 7: Backup path success (asymmetric mode) in contrast to possible single-indirect solutions and randomized approach.

Similarly, Fig. 8 depicts the results concerning the symmetric mode. The prior findings also apply to this mode, even if the symmetric mode performs slightly inferior to the asymmetric one. This performance advantage of the asymmetric mode, however, comes at the cost of few additional VPN connections.

5.2 Impact on the VPN node degree

The overlay node degree was identified as major criterion in a sense of scalability and efficiency. Typically, the reference systems nodes maintain a low, equally distributed number of overlay connections proactively. To estimate the increased node degree, the proposed backup path system is compared to an unextended environment within a VPN scenario containing 100 nodes.

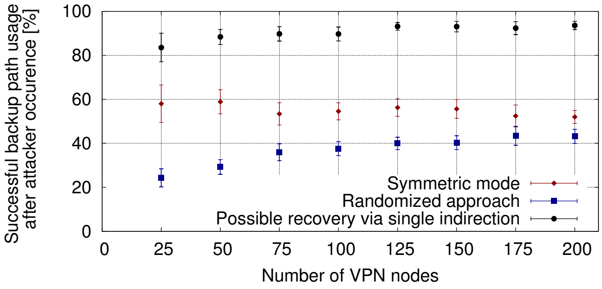

Figure 8: Backup path success (symmetric mode) in contrast to possible single-indirect solutions and randomized approach

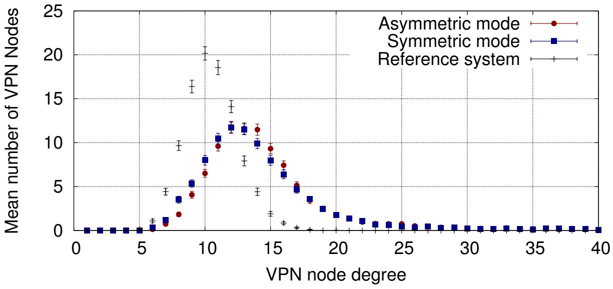

Figure 9: Node degree distribution while using symmetric or asymmetric backup path mode.

Fig. 9 shows the resulting node degree distribution for the reference system as well as for the symmetric and asymmetric mode. The node distribution of the native reference system environment gives a vague idea of a normal distribution form around mean 10.5, whereas the backup paths mechanism results in a shift to 14.5 in symmetric and 14.9 in asymmetric mode. Also the empirical standard deviation increases from 2.0 for reference system up to 5.9 for the asymmetric mode, which is due to the higher number of VPN nodes with large degrees. This behavior can be attributed to the preferred selection of some more attractive backup mediation nodes. Especially nodes within the same AS make similar choices regarding their backup paths towards destination peers and due to the exponentially distributed size of VPN groups the chosen mediators will suffer from higher node degrees. However, this is not considered to pose a practical problem for most scenarios, as potential mediators may reject backup requests, if burdened too much. Furthermore, due to the proactive creation of the paths, the mediation can take place in idle times and does not affect the performance of undisturbed VPNs.

6. CONCLUSION

Even though attacks on the BGP routing are currently a mostly theoretical problem, their effectiveness will make them a major threat for VPNs. However, the presented backup path mechanism will allow dynamic VPNs to proactively prepare for such situations at the cost of a comparably low overhead. Due to the use of local knowledge and the LCS metric, the system is also robust against attacks that target at a circumvention of the backup path mechanism itself. For future research it is planned to integrate the topology measurements in a metric, which also rates factors like jitter, drop rate, history of the path, and detected load balancers. This allows for a better reaction to effects of dynamic attacks. Here, a major challenge is the creation of realistic evaluation scenarios to determine the robustness against such attacks. Furthermore, we want to evaluate the effectiveness of our approach in terms of prevention of overlay partitions, as the creation of topological divergent paths makes it more difficult for attackers to split VPNs.

7. REFERENCES

[1] The Cooperative Association for Internet Data Analysis (CAIDA), 2011.

[2] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris. Resilient Overlay Networks. In *SOSP '01: Proceedings of the eighteenth ACM symposium on Operating systems principles*, pages 131–145, 2001.

[3] D. Andersen, A. Snoeren, and H. Balakrishnan. Best-path vs. multi-path overlay routing. In *Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement*, pages 91–100. ACM, 2003.

[4] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris. The Case for Resilient Overlay Networks. In *Proc. of the 8th Annual Workshop on Hot Topics in Operating Systems*, pages 152–157, 2001.

[5] A. Collins. *The Detour Framework for Packet Rerouting.* PhD thesis, University of Washington, 1998.

[6] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to algorithms*, chapter 15.4 Longest Common Subsequence, pages 390–397. The MIT press, 3rd edition, 2009.

[7] J. Cowie, A. Ogielski, B. Premore, and Y. Yuan. Internet worms and global routing instabilities. In *Proceedings of SPIE*, volume 4868, page 195, 2002.

[8] D. Danchev. Coordinated Russia vs Georgia cyber attack in progress, Aug. 2008.

[9] Forrester Consulting. The Trends And Changing Landscape Of DDoS Threats And Protection, July 2009.

[10] P. B. Gentry. What is a VPN? *Information Security Technical Report*, 6(1):15–22, 2001.

[11] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian. Delayed Internet routing convergence. *ACM SIGCOMM Computer Communication Review*, 30(4):175–187, 2000.

[12] V. Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady*, 10:707–710, 1966.

[13] S. Mansfield-Devine. Anonymous: serious threat or mere annoyance? *Network Security*, 2011(1):4–10, 2011.

[14] R. Musunuri and J. A. Cobb. An overview of solutions to avoid persistent BGP divergence. *IEEE Network*, 19(6):28–34, 2005.

[15] A. Nakao, L. Peterson, and A. Bavier. A routing underlay for overlay networks. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 11–18. ACM, 2003.

[16] V. Paxson. End-to-end routing behavior in the internet. *SIGCOMM Comput. Commun. Rev.*, 26(4):25–38, August 1996.

[17] V. Paxson. End-to-end internet packet dynamics. *SIGCOMM Comput. Commun. Rev.*, 27(4):139–152, October 1997.

[18] M. Rossberg, G. Schaefer, and T. Strufe. Distributed Automatic Configuration of Complex IPsec-Infrastructures. *Journal of Network and Systems Management*, 18(3):300–326, 2010.

[19] S. Savage, T. Anderson, A. Aggarwal, D. Becker, N. Cardwell, A. Collins, E. Hoffman, J. Snell, A. Vahdat, G. Voelker, et al. Detour: Informed Internet Routing and Transport. *IEEE Micro*, 19(1):50–59, 1999.

[20] S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson. The End-to-End Effects of Internet Path Selection. *Computer Communication Review*, 29(4):289–299, 1999.

[21] D. Sontag, Y. Zhang, A. Phanishayee, D. G. Andersen, and D. Karger. Scaling All-Pairs Overlay Routing. In *Proceedings of the 5th international conference on Emerging networking experiments and technologies (CoNEXT)*, pages 145–156, 2009.