

Rushing Attacks and Defense in Wireless Ad Hoc Network Routing Protocols

Yih-Chun Hu, Adrian Perrig & David B. Johnson



BGP Misconfiguration

- It is well-known that simple, accidental BGP configuration errors can disrupt Internet connectivity.
- Yet little is known about the frequency of misconfiguration or its causes, except for the few spectacular incidents of widespread outages.
 - In this paper, we present the first quantitative study of BGP misconfiguration.
 - Over a three week period, we analyzed routing table advertisements from 23 vantage points across the Internet backbone to detect incidents of misconfiguration.
 - For each incident we polled the ISP operators involved to verify whether it was a misconfiguration, and to learn the cause of the incident.



Misconfiguring BGP

- We also actively probed the Internet to determine the impact of misconfiguration on connectivity.
- Surprisingly, we find that configuration errors are pervasive, with 200-1200 prefixes (0.2-1.0% of the BGP table size) suffering from misconfiguration each day.
- Close to 3 in 4 of all new prefix advertisements were results of misconfiguration.
- Fortunately, the connectivity seen by end users is surprisingly robust to misconfigurations.
- While misconfigurations can substantially increase router processing overhead, only one in twenty five affects connectivity.
- While the causes of misconfiguration are diverse, we argue that most could be prevented through better router design.



What is BGP? RFC 1771, 1772

- In the Internet environment. BGP is an inter-Autonomous System routing protocol. The network reachability information exchanged via BGP provides sufficient information to detect routing loops and enforce routing decisions based on performance preference and policy constraints as outlined in [RFC 1104](#)
- In particular, BGP exchanges routing information containing full AS paths and enforces routing policies based on configuration information.

Factors driving this RFC:

1. Exhaustion of the class-B network address space. One fundamental cause of this problem is the lack of a network class of a size which is appropriate for mid-sized organization; class-C, with a maximum of 254 host addresses, is too small while class-B, which allows up to 65534 addresses, is too large to be densely populated.
2. Growth of routing tables in Internet routers are beyond the ability of current software (and people) to effectively manage.
3. Eventual exhaustion of the 32-bit IP address space.



Classless inter-domain routing (CIDR)

- **Classless inter-domain routing (CIDR) attempts to deal with these problems by proposing a mechanism to slow the growth of the routing table and the need for allocating new IP network numbers.**
- **It does not attempt to solve the third problem, which is of a more long-term nature, but instead endeavors to ease enough of the short to mid-term difficulties to allow the Internet to continue to function efficiently while progress is made on a longer-term solution.**
- **BGP-4 is an extension of BGP-3 that provides support for routing information aggregation and reduction based on the Classless inter-domain routing architecture (CIDR)**



CIDR Architecture RFC 1519

- The proposed solution is to topologically allocate future IP address assignment, by allocating segments of the IP address space to the transit routing domains.
- There are two basic components of this addressing and routing plan: one, to distribute the allocation of Internet address space and two, to provide a mechanism for the aggregation of routing information.



Aggregation Routing

- Aggregation and its limitations
 - One major goal of the CIDR addressing plan is to allocate Internet address space in such a manner as to allow aggregation of routing information along topological lines.
 - For simple, single-homed clients, the allocation of their address space out of a transit routing domain's space will accomplish this automatically - rather than advertise a separate route for each such client, the transit domain may advertise a single aggregate route which describes all of the destinations connected to it.
 - Unfortunately, not all sites are singly-connected to the network, so some loss of ability to aggregate is realized for the non-trivial cases.



Autonomous Systems

- All of the discussions in RFC 1772 are based on the assumption that the Internet is a collection of arbitrarily connected Autonomous Systems.
- That is, the Internet will be modeled as a general graph whose nodes are AS's and whose edges are connections between pairs of AS's.
- The classic definition of an Autonomous System is a set of routers under a single technical administration, using an interior gateway protocol and common metrics to route packets within the AS and using an exterior gateway protocol to route packets to other AS's.
- Since this classic definition was developed, it has become common for a single AS to use several interior gateway protocols and sometimes several sets of metrics within an AS.
- The use of the term Autonomous System here stresses the fact that, even when multiple IGPs and metrics are used, the administration of an AS appears to other AS's to have a single coherent interior routing plan and presents a consistent picture of which destinations are reachable through it.
- AS's are assumed to be administered by a single administrative entity, at least for the purposes of representation of routing information to systems outside of the AS.



BGP Topological Model

- When we say that a connection exists between two AS's, we mean two things:
- **Physical connection:** There is a shared Data Link subnetwork between the two AS's, and on this shared subnetwork each AS has at least one border gateway belonging to that AS. Thus the border gateway of each AS can forward packets to the border gateway of the other AS without resorting to Inter-AS or Intra-AS routing.
- **BGP connection:** There is a BGP session between BGP speakers in each of the AS's, and this session communicates those routes that can be used for specific destinations via the advertising AS.



On-Demand Protocols

- There are two categories of routing protocols: table-driven and on demand-routing.
- In table-driven routing protocols routing information is periodically advertised to all nodes so all nodes have an up-to-date view of the network.
- Alternatively, on-demand routing protocols only discovers a new route when it is required to.
- Hybrid routing protocols also exist and they try to achieve an efficient balance between both categories of protocols



Comparison between Table-Drive Routing and On-Demand Routing

	Table-driven Routing	On-Demand Routing
Availability of Routing Information	Immediately from route table	After a route discovery
Route Updates	Periodic Advertisements	When requested
Routing Overhead	Proportional to the size of the network regardless of network traffic	Proportional to the number of communicating nodes and increases with increased node mobility

It is clear that on-demand protocols are more suited for mobile handheld devices as network bandwidth and battery power is limited.



Ad hoc On-demand Distance Vector Routing (AODV)

- Ad hoc On-demand Distance Vector Routing (AODV) is an on-demand version of the table-driven Dynamic Destination-Sequenced Distance-Vector (DSDV) protocol
- To find a route to the destination, the source broadcasts a route request packet.
- This broadcast message propagates through the network until it reaches an intermediate node that has recent route information about the destination or until it reaches the destination.
- When intermediate nodes forwards the route request packet it records in its own tables which node the route request came from.
- This information is used to form the reply path for the route reply packet as AODV uses only symmetric links.
- As the route reply packet traverses back to the source, the nodes along the reverse path enter the routing information into their tables.
- When ever a link failure occurs, the source is notified and a route discovery can be requested again if needed.



Dynamic Source Routing

- The Dynamic Source Routing (DSR) protocol is a source-routed on-demand protocol.
- There are two major phases for the protocol: route discovery and route maintenance.
- The key difference between DSR and other protocols is the routing information is contained in the packet header.
- Since the routing information is contained in the packet header then the intermediate nodes do not need to maintain routing information.
- An intermediate node may wish to record the routing information in its tables to improve performance but it is not mandatory.
- Another feature of DSR is that it supports asymmetric links as a route reply can be piggybacked onto a new route request packet.
- DSR is suited for small to medium sized networks as its overhead can scale all the way down to zero.
- The overhead will increase significantly for networks with larger hop diameters as more routing information will be contained in the packet headers



Rushing attack

- Disseminates ROUTE REQUESTs quickly throughout the network, suppressing any later legitimate ROUTE REQUESTs when nodes drop them due to the duplicate suppression.
- Practical only on on-demand routing protocols using duplicate suppression at each node (AODV)
- Threats: Failure of route discovery



Rushing Attack Example




Figure 1: Example network illustrating the rushing attack.

- The initiator node initiates a Route Discovery for the target node.
- If the Route Requests for this Discovery forwarded by the attacker are the first to reach each neighbor of the target, then any route discovered by this Route Discovery will include a hop through the attacker.

Assumptions

- We make the common assumption that most network links are bidirectional.
- Though a jamming attack is also an important denial-of-service attack, we present mechanisms to defend against the rushing attack because we believe that the rushing attack is more easily performed.
- Medium Access Control protocols are also often vulnerable to attack.
 - For example, in IEEE 802.11, an attacker can paralyze nodes in its neighborhood by sending Clear-To-Send (CTS) frames periodically, setting the “Duration” field of each frame equal to the interval between such frames.
- Prior work has shown that ad hoc network routing in general does not scale well.
 - Most existing simulation of ad hoc network routing protocols consider scenarios of 50 to 500 nodes.



SECURE ROUTING REQUIREMENTS AND PROTOCOL

- In this section, we describe a set of generic mechanisms that together defend against the rushing attack: *secure Neighbor Detection*, *secure route delegation*, and *randomized ROUTE REQUEST forwarding*.

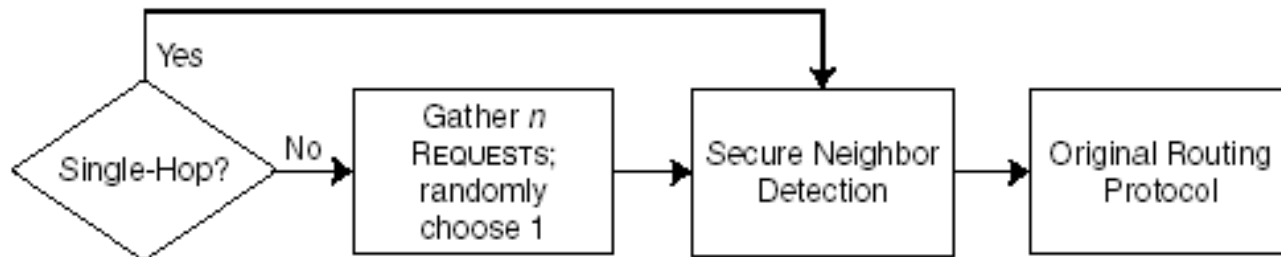


Figure 2: Our combined mechanisms to secure an on-demand route discovery protocol against the rushing attack.



Secure Neighbor Detection

S : $\eta_1 \xleftarrow{R} \{0, 1\}^\ell$
 $M_1 = \langle \text{NEIGHBOR SOLICITATION}, S, \eta_1 \rangle$
 $\Sigma_{M_1} = \text{Sign}(H(M_1))$

$S \rightarrow * :$ $\langle M_1, \Sigma_{M_1} \rangle$

R : $\eta_2 \xleftarrow{R} \{0, 1\}^\ell$
 $M_2 = \langle \text{NEIGHBOR REPLY}, S, R, \eta_1, \eta_2 \rangle$
 $\Sigma_{M_2} = \text{Sign}(H(M_2))$

$R \rightarrow S :$ $\langle M_2, \Sigma_{M_2} \rangle$

S : $M_3 = \langle \text{NEIGHBOR VERIFICATION}, S, R, \eta_1, \eta_2 \rangle$
 $\Sigma_{M_3} = \text{Sign}(H(M_3))$

$S \rightarrow R :$ $\langle M_3, \Sigma_{M_3} \rangle$

Figure 3: Neighbor Detection between initiator S and responder R .



Secure Route Delegation

- **S-BGP uses Route Attestations to ensure that each AS listed in the BGP AS path is indeed a valid AS**
- **In S-BGP, before sending a route update to its neighbor, the AS signs a route attestation delegating it the right to further propagate the update.**
- **Used to make sure that BOTH neighbors believe that they are within transmission range.**



Randomized Message Forwarding

- The secure Neighbor Detection and secure Route Delegation techniques are not sufficient to thwart the rushing attack, since an adversary can still get an advantage by forwarding ROUTE REQUESTs very rapidly.
- We use a random selection technique to minimize the chance that a rushing adversary can dominate all returned routes.
- In traditional ROUTE REQUEST forwarding, the receiving node immediately forwards the REQUEST and suppresses all subsequent REQUESTs.
- In our modified flooding, a node first collects a number of REQUESTs, and selects a REQUEST at random to forward.
- There are thus two parameters to our randomized forwarding technique: first, the number of REQUEST packets to be collected, and second, the algorithm by which timeouts are chosen.
- Given perfect information, each forwarding node would collect the maximum possible number of REQUESTs before forwarding one, since this approach provides the most effective defense against a rushing attack.



Randomized Message Forwarding (2)

- However, when the number of REQUESTs is chosen to be too large, randomized forwarding will heavily rely on the timeout to trigger REQUEST forwarding, increasing latency and possibly reducing security.
- In a real network, perfect information is generally not available; as a result, initiators can include in each Route Discovery the number of REQUESTs to buffer before forwarding one, and can adjust this parameter adaptively, based on the REPLY latency and on the parameters chosen by other nodes.
 - Alternatively, this number can be chosen as a global parameter, or locally using an adaptive algorithm, though an adaptive algorithm may allow certain new attacks.
- With topology information, the choice of timeout should be based on the number of legitimate hops between the initiator and the node forwarding the REQUEST; closer nodes should choose shorter timeouts than far-away nodes.
 - This topological information can be approximated by location information; that is, nodes that are geographically closer should choose smaller timeouts than nodes that are geographically farther away.
- When geographic information is not available, nodes can randomly choose timeouts; however, this approach reduces security by favoring nodes choosing shorter timeouts.



Randomized Route Request Forwarding (Secure Route Discovery)

- The intuition behind Secure Route Discovery is to make the forwarding of REQUEST packets less predictable by buffering the first n REQUESTs received, then randomly choosing one of those REQUESTs.
- However, we need to prevent an attacker from filling too many of these n REQUESTs, since otherwise the attacker could simply rush n copies of a REQUEST, rather than a single one. Our scheme would once again be vulnerable to the rushing attack.




Secure Route Discovery


- To limit the number of REQUESTs that traverse an attacker, we exploit the fact that legitimate nodes forward only one REQUEST in any Discovery.
 - First, we require that each REQUEST carry a list of nodes traversed by this REQUEST.
 - Second, we require a bidirectional Neighbor Verification for each link represented by this list of nodes, for a total of two signed Neighbor Verifications per hop.
 - Third, to authenticate the node list, we require each node to authenticate the REQUEST it forwards, though it can piggyback this authentication together with the Neighbor Verification that it signs.
 - Finally, we require buffered REQUESTs be duplicate-suppression-unique:
 - that is, if the route record of any two REQUESTs contain any node A , the route prefix leading up to (and including) A must be the same. These three requirements constrain an attacker to the extent that an attacker that has compromised m nodes can rush at most m REQUESTs.
- To prevent replay of old Neighbor Verification messages, each message is tied to a specific Route Discovery.




Simulation Evaluations




(a) Packet Delivery Ratio



(b) Median Latency



(c) Packet Overhead



(d) Byte Overhead

Figure 5: Unoptimized RAP performance evaluation results in non-adversarial environment. Optimized RAP would have same results as Ariadne, except that it would perform better when under attack. Under attack, optimized RAP and Ariadne would perform identically for one- and two-hop routes, but in finding longer routes, RAP should significantly outperform Ariadne, since RAP finds working routes with moderate probability, but Ariadne and DSR can never find routes. “RAP 1 Flow” refers to RAP with the lighter communications pattern of one CBR source. Results based on averages over 50 simulation runs; the error bars represent the 95% confidence interval of the mean.



Conclusions

- In this paper, we have described the *rushing attack*, a novel and powerful attack against on-demand ad hoc network routing protocols.
 - This attack allows an attacker to mount a denial-of-service attack against *all* previously proposed secure on-demand ad hoc network routing protocols.
 - We have also presented RAP (Rushing Attack Prevention), a new protocol that thwarts the rushing attack.
- We found that the widely used duplicate suppression technique makes the rushing attack possible, and we designed a new Route Discovery protocol called RAP that replaces the standard mechanism and thwarts the rushing attack.
 - Our approach is generic, so any protocol that relies on duplicate suppression in Route Discovery can use our results to fend off rushing attacks.
- More importantly, we demonstrated that there are mechanisms that can defend against the rushing attack, even though all previous attempts at secure on-demand ad hoc network routing protocols have been vulnerable.
- When integrated with a secure routing protocol, RAP incurs *no cost* unless the underlying secure protocol cannot find valid routes.
- When RAP is enabled, it incurs higher overhead than do standard in Link State Routing.

