BGP SECURITY

TCOM 610

George Mason University

Tawfiq Khan

BGP Vulnerability

- BGP is built on top of TCP: BGP implementation is vulnerable to all TCP protocol vulnerabilities (Session hijack, SYN flood etc)
- BGP update messages contain routing control info: malicious BGP updates can potentially disrupt Internet routing operations
- BGP process is run on router/server with limited CPU/Memory and this process is vulnerable to DOS attacks

BGP Spoofing Attacks

- BGP Spoofing attacks are those in which the BGP peer is imitated
  - Can be TCP based spoofing targeting the BGP port of the router or spoofed BGP packets
- To successfully spoof a TCP session supporting the BGP peers, the following must be achieved:
  - Source IP address must be spoofed
  - Source Port must be spoofed
  - TCP Sequence Number must match
  - IP's TTL must match
  - Destination port must match (not always 179, depending on which side initiates the communication)

Spoofing Countermeasures

- MD5 on BGP peering session mitigates most wire sniffing threats to BGP Spoofing
  - Adding password-based message digest makes those spoofed packets automatically drop
- Use of diverse keys on eBGP session with fellow ISPs would mitigate the risk of the MD5 key from leaking
- If operationally feasible, treating MD5 keys with changes policies similar to password change policies
- Difficulties with MD5 key maintenance within an operational ISP environment
- MD5 has been more widely implemented due to some recent BGP implementation vulnerabilities

BGP Hijacking

- Requires a successful BGP Spoof
- Attacks masquerade BGP status packets as coming from the neighbor. The packets would look legitimate, but would carry malicious BGP status updates.
- The updates could be tearing down the BGP session, inserting routing information, or withdrawing valid routing information.
- Effective BGP Hijacking requires additional knowledge of the current BGP interaction between the peers

Hijacking Countermeasures

- MD5 on BGP peering session mitigates most wire sniffing threats to BGP Hijacking.
- Work is in progress on a BGP over IPSEC option or sBGP that would greatly increase the difficulty of hijacking
- There is big operational burden for deploying BGP over IPSec or sBGP
- The current consensus is that we may not need security

De-aggregation Attacks

- Announcement of more specific routes (/24s) for practically the entire Internet
- Due to the longest match routing policy, the most specific route will always win
- Will consume router memory, disrupt global Internet routing operation and may even crash routers
- In some cases, saturated links cause more outages
- Multi-homed customers with BGP speaking routers could be broken into and used to launch a de-aggregation attack

De-aggregation Attacks

- Max Prefix Limits on peer connections combined with aggressive route filtering of the ISP's customers effectively mitigates the de-aggregation risk.
- ISPs should only permit customer prefixes for those IP address blocks that have been allocated to them by the IANA system.
- These IP allocation records can be validated through the RIR databases, their customers, and their peers (if the customer is a multi-homed customer)
- ISP route filtering for not accepting more specific routes

Un-authorized Route Injection Attacks

- Advertisement of routes in which the network does not have allocation authority pulls traffic away from the authorized network.
- This causes a DOS on the network who allocated the block of addresses (no traffic) and may cause a DOS on the network in which it re-advertised.
- The risk increases as more enterprise are multi-homed connected
- The easiest attack vector being advertisement of someone else's IP address block

Countermeasure to Un-authorized Route Injection Attacks

- Aggressive egress routing filtering prefixes set to other ISPs on the peering points (and customers) mitigate the risk of malicious advertisement of un-authorized routes into the Global Internet Route table.
- With significant limitations, these IP allocation records can be validated through the RIR databases, the ISP's customer databases, and their peer contacts (if the customer is a multi-homed customer) – not automatically and dynamically
- This egress filtering contains malicious advertisement from a violated router within an ISP's Autonomous System – keeping the advertisement from spreading to other ISPs

Unallocated Route Injection Attack

- Advertisement of IP addresses that have yet to be allocated by IANA can pose several problems on the Internet:
  - BGP table explosions
  - The use of latent backscatter as a DOS tool
- Most ISPs do not filter Bogons (packets with unauthorized routes) – the term used to describe the IANA reserved address space
- Malicious attack might use a violated BGP speaking router to start advertising large ranges of Bogon space
  - Overloading BGP and forwarding tables in routers
  - Turn the advertising router into an *Internet Sink Hole* Bogon Route
- Filtering – filtering all address blocks that have yet to be allocated – is an effective counter to this attack vector.
- IANA maintains public list of Ipv4 allocations. The *IANA Reserved* blocks are the Bogon blocks.

Resource Saturation Attack

- DOS/DDOS Attacks directly against the BGP protocol port (port 179) are perceived to be an easily executable attack vector
- TCP syn flood against port 179 attempt to flood the application port.
- Actually, they end up flooding a resource like the input queue, forcing the router's processors to work over time with queue maintenance. At times, queue and processor resources can reach the point where control plane packets are dropped. When control plane traffic is dropped, the routing protocol sessions drop resulting in a route flap

DOS Attack Trend

- What's the trend in attacks ?
  - Yesterday: bandwidth abuse, exploiting bugs
  - Today: packets-per-second, also against (core) routers
  - Tomorrow:
    - QoS/"extended" header
    - (InterAS) MPLS VPNs' trust model
    - IPv6 (transition)
    - Somewhere in the forwarding path code
    - Non-spoofed sources (who cares if you have 100k+ bots anyway)
    - Protocol complexity attacks (mixed with/hidden in/part of "normal" traffic): ie. low bandwidth "special" packets
    - Is the issue really BGP/DNS hijacking ?

Resource Saturation Attack Countermeasures

- ACLs to protect BGP mitigate some direct attacks, but not spoofed attacks
  - Spoofed attacks only need to match the source IP address of the BGP peer; once spoofed, the packet passes right through the ACL
- IP Source validation on the edge of an ISP's network would also help mitigate the risk,
- Proposed BGP mitigation techniques are just as vulnerable to these sorts of *saturation attacks*. BGP with MD5, BGP over IPSEC, or other BGP security proposals are all exposed to resource saturation attacks
- Increase the input queue depth to the point where router has enough room to drop the attack packets and still have room to keep the control plane traffic
- TCP state management techniques that would not respond or clear out SYN and SYN/ACK floods
- Multi-layered/multi-level redundancy designed used on today's ISP networks allow for the back-up path to maintain network integrity

BGP MD5 Authentications (RFC2385)

- Design to protect BGP session authenticity: against the introduction of spoofed TCP segments into the connection stream, especially RST packet (needs to match Sequence number)
- All TCP packets will contain16-byte MD5 digest produced by applying the MD5 algorithm to:
  - TCP pseudo-header (in the order: source IP address, destination IP address, zero-padded protocol number, and segment length)
  - TCP header, excluding options, and assuming a checksum of zero
  - TCP segment data (if any)
  - an independently-specified key or password, known to both sides and presumably connection-specific

MD5 Message Format

The proposed option has the following format:

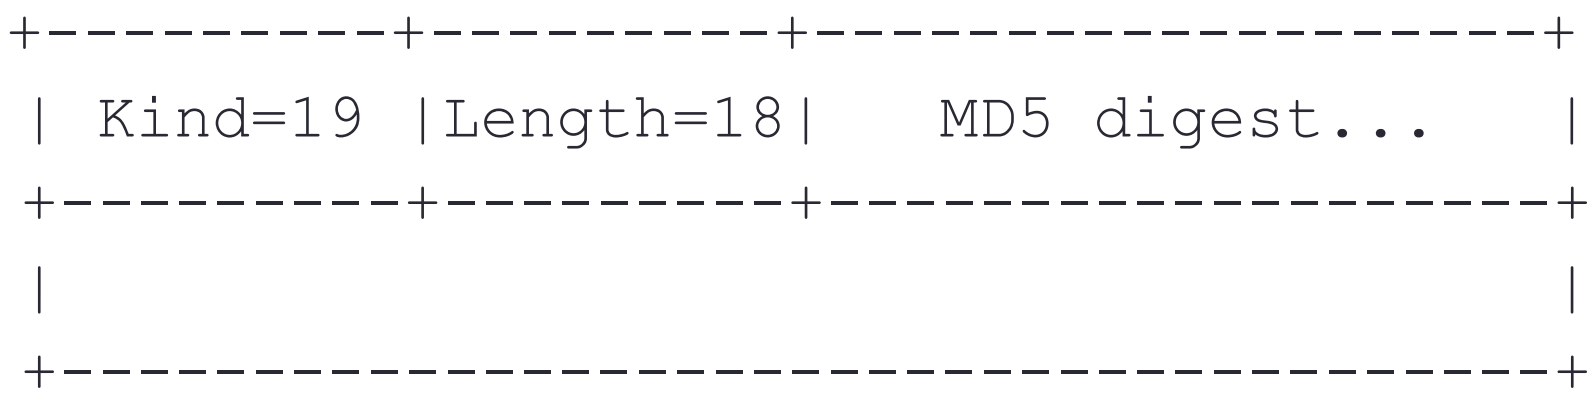

The MD5 digest is always 16 bytes in length.

MD5 Operations

- Connectionless reset: RST packet will be ignored since the originator does not have the key to generate proper signature for the segment
- Performance: calculating and verifying each TCP segment (both inbound and outbound) requires CPU
- As of today, most of the large ISP's EBGP session has been enhanced to enable MD5
- For several recent BGP vulnerability, MD5 is the recommended workaround

MD5 In Action

- On April 20, 2004: BGP MD5 makes into headline due to new vulnerability found in most router vendor's BGP implementation
- Initial Vulnerability report: http://www.uniras.gov.uk/vuls/2004/236929/index.htm
- RFC for the potential fix: http://www.ietf.org/internet-drafts/draft-ietf-tcpm-tcpsecure-00.txt
- As the result of this vulnerability, all eBGP sessions for most ISP have been updated to MD5 authentication
- MD5 key management and update is an operation challenge

BGP Route Authenticity

- If an established BGP peer sends you updates with new routes, how can you verify them or trust them?
- Lack of a scalable means of verifying the authenticity and legitimacy of BGP control traffic (sBGP)
- Lack of the authentication for the origin of any advertisement within BGP can be verified and authenticated, and verification that the final destination in the path is actually within peer's AS (soBGP)
- The use of a packet's Time to Live (TTL) (IPv4) or Hop Limit (IPv6) to protect a protocol stack from CPU-utilization based attacks (RFC3682)

sBGP

- Address-based PKI is used to validate signatures
  - Authentication of ownership for IP address blocks, AS number, an AS's identity, and a BGP router's identity
  - Use existing infrastructure (Internet registries etc.)
  - Routing origination is digitally signed
  - BGP updates are digitally signed
- A new, optional, BGP transitive path attribute is employed to carry digital signatures covering the routing information in a BGP UPDATE
  - Receivers can verify the address prefixes and path information
- IPsec is used to provide data and partial sequence integrity, and to enable BGP routers to authenticate each other for exchanges of BGP control traffic
- Pre-distribute (most) certificates to near each BGP speaker; Cache signed routes and originations

sBGP

- Costs
  - Bandwidth overhead
  - CPU – lots of CPU, may need hardware assistance
  - Memory for additional transitive attributes
  - Setting up PKI
- Limitations:
  - PKI is complex
  - Does not authenticate route withdrawal
  - Requires router upgrade
- More Info: http://www.ir.bbn.com/projects/s-bgp/
- Latest RFC draft: draft-clynn-s-bgp-protocol-01.txt – June, 2003

soBGP

- **soBGP (Secure Origin BGP) targets the need to verify the validity of an advertised prefix**
- **Goal: Validate that an AS is authorized to originate a prefix**
- **Design requirements**
  - **Take advantage of existing Internet routing operational experience**
  - **Minimize impact to current BGP implementations**
  - **Must *not* rely on a central authority of any type**
  - **Should not rely on routing to secure routing**
  - **Must be incrementally deployable**
  - **Must allow deployment flexibility**
  - **Flexibility should be provided to allow operators to configure the level of security vs. overhead and convergence speed**

soBGP

- New BGP SECURITY Message used to carry security information
  - Certificates are carried within TLVs
  - Expandable to other security related information
  - Negotiated at session startup (capability exchange)
- Verifies that the originator of a route is authorized to do so
  - Verifies that the advertised AS_PATH represents a valid path to the originator
- Fixed additional scalability requirements
  - Per-AS information and route policies advertised once.
  - No additional information in UPDATES, resulting in low processing impact.
- Use of Certificates to advertise and correlate AS identity, prefix ownership and route policy
  - Entity Certificate = Used to establish identity
  - Authorization Certificate = Used to assign and delegate IP address space
  - Policy Certificate = Used to define per-AS or pre-prefix
  - policies and propagate AS interconnectivity topology map

soBGP

- **Uses Web-of-Trust model to validate certificates.**
- **No specific root (single point of failure), but distributed responsibility**
- **Built in Flexibility**
  - **UPDATE and Certificate propagation may be decoupled.**
  - **On or off-box cryptography operations (inside the local AS)**
  - **Incrementally deployable – provides some security in any multi-AS scenario.**
  - **Configurable level of validation and weights**

soBGP Deployment

- **Scenario One:**
  - **Exchange certificates at all eBGP peering points (AS edges).**
  - **Process the certificates, and build the required soBGP tables at each eBGP speaker**
- **Scenarios Two:**
  - **Certificates can also be exchanged at the AS edge, and “shuttled,” using iBGP connections, to a server within the AS.**
  - **These servers then perform all certificate processing, and build the necessary databases.**
  - **The edge routers then consult these servers, using RADIUS, to validate received updates.**
- **Other scenarios:**
  - **Certificates can also be exchanged, using multihop eBGP directly between the soBGP servers in each AS**
  - **Certificates may also be exchanged with third party providers of some type**
  - **Certificates may be generated by one AS, and advertised by another AS**

soBGP Update

- **ftp://ftp-eng.cisco.com/sobgp**
- Latest drafts:
  - **draft-ng-sobgp-bgp-extensions –xx**
  - **draft-white-sobgp-bgp-deployment- xx**

BGP TTL Hack (RFC3682)

- Generalized TTL Security Mechanism (GTSM) is designed to protect a router's TCP/IP based control plane from CPU-utilization based attacks
- Assumptions:
  - Vast majority of protocol peerings are established between routers that are adjacent
  - TTL spoofing is considered nearly impossible
- GTSM mechanism is equally applicable to both TTL (IPv4) and Hop Limit (IPv6)
- It is common practice for many service providers to ingress filter (deny) packets that have the provider's loopback addresses as the source IP address
- The router supports a method of classifying traffic destined for the route processor into interesting/control and not-control queues

GTSM

- For directly connected routers, Set the outbound TTL for the protocol connection to 255.
- For each configured protocol peer:
  - Update the receive path Access Control List (ACL) or firewall to only allow protocol packets to pass onto the Route Processor (RP) that have the correct <source, destination, TTL> tuple
  - The TTL must either be 255 (for a directly connected peer), or 255-(configured- range-of-acceptable-hops) for a multi-hop peer.
  - It is assumed that a receive path ACL is an ACL that is designed to control which packets are allowed to go to the RP

GTSM Limitations

- Hard to deal with multi-hop scenarios: TTL may change due to routing events
- Not applicable to IBGP
- Need to deal with tunneling and MPLS TTL behaviors

Cisco CRS Features

- Control Plane Protection
  - Distributed and redundant route processor
  - MD5, ACL filter and QoS rate-limiting
  - Dynamic Control Plane Protection (DCPP): configured BGP peers automatically allocated adequate resources
  - Automatic Control Plane Congestion Filter
  - GTSM – TTL
  - RPL: Routing policy language – modular and hierarchical policy language

IETF RPSEC Group

- Improper route origination or propagation
  - Accidental and naive
    - 18.0.0.0/8
    - AS 7007
  - Malicious
- Antagonistic or competing announcements
- "Pop-up" hacking using un-announced space
  - advertise, hack or spam, disappear
- Is route filtering really the solution?
  - How do we check who's authorized, really?

IETF RPSEC

- Session resets, data injection/corruption
  - Blind attacks
- TCP-MD5: solution or other problems?
  - Key management (How to pass keys around?)
  - Rogue employees (Who wants some keys?)
  - Compromised routers (Who needs keys?)
- GTSH (Generalized TTL Security Hack)?
  - As above for bad routers/employees

IETF BGP Security

- **The RPsec WG within the IETF is currently documenting BGP security requirements.**
- **In Scope:**
  - Originating False Data, Routing Database
  - Integrity, Peering Integrity
- **Out of Scope:**
  - Any attack where BGP isn’t directly
  - manipulated, Data Packet Delivery

IETF BGP Security

- **MUST Support A Distributed Trust Model**
  - The optimal trust model may vary.
  - A strict hierarchy is a subset of a distributed trust model.
- **Routing Information Validation**
  - **MUST Verify Origin AS' Authorization to Advertise**
  - The trust model is critical in this area.
- **MUST verify the AS path corresponds to a valid path in the Internetworking**
- **MUST ensure the first element of the AS path is the same as the transmitting peer's AS**
- **Logging/Tracking**
  - **SHOULD Provide Non-Repudiation of Updates**
  - **MUST Provide for Logging**
- **MUST Include Transport Protection Between BGP Speakers**