DE-CIX blackholing service

Meeting of the eco/DE-CIX competence group security

Sebastian Abt

da/sec – Biometrics and Internet Security Research Group
Hochschule Darmstadt, Darmstadt, Germany

Frankfurt a.M., November 14th, 2012

Why do we meet today?

- DE-CIX announced a new blackholing service at its customer summit on August 28th, 2012.
- A longer discussion about pitfalls related to this service started at tech-list on September 19th, 2012 (thread length: 20; longest thread in 2012).

$\Rightarrow$ The topic is relevant and has possible security impact (availability). We should discuss it in this group and draw appropriate conclusions.

Why do we meet today?

- DE-CIX announced a new blackholing service at its customer summit on August 28th, 2012.
- A longer discussion about pitfalls related to this service started at tech-list on September 19th, 2012 (thread length: 20; longest thread in 2012).

$\Rightarrow$ The topic is relevant and has possible security impact (availability). We should discuss it in this group and draw appropriate conclusions.

Possible outcome of this meeting

- Common understanding of opportunities and pitfalls associated with DE-CIX blackholing service.
- Recommendation for service implementation.
- Best practices for service utilization.
- Examples of BGP peer configurations.

- **Question:** Does anybody know if a similar service is offered at any other IXP?

Possible outcome of this meeting

- Common understanding of opportunities and pitfalls associated with DE-CIX blackholing service.
- Recommendation for service implementation.
- Best practices for service utilization.
- Examples of BGP peer configurations.

- **Question:** Does anybody know if a similar service is offered at any other IXP?

Goal of DE-CIX blackholing service

Enable ISP V under attack to easily and effectively block incoming (attack) traffic destined towards prefix p at IXP level before it hits ISP V's infrastructure.

- Transfer of well established concept between, e.g., upstream and downstream.
- However, we should remark that this is just the last line of defense.

Goal of DE-CIX blackholing service

Enable ISP V under attack to easily and effectively block incoming (attack) traffic destined towards prefix p at IXP level before it hits ISP V's infrastructure.

- Transfer of well established concept between, e.g., upstream and downstream.
- However, we should remark that this is just the last line of defense.

Current implementation of blackholing service

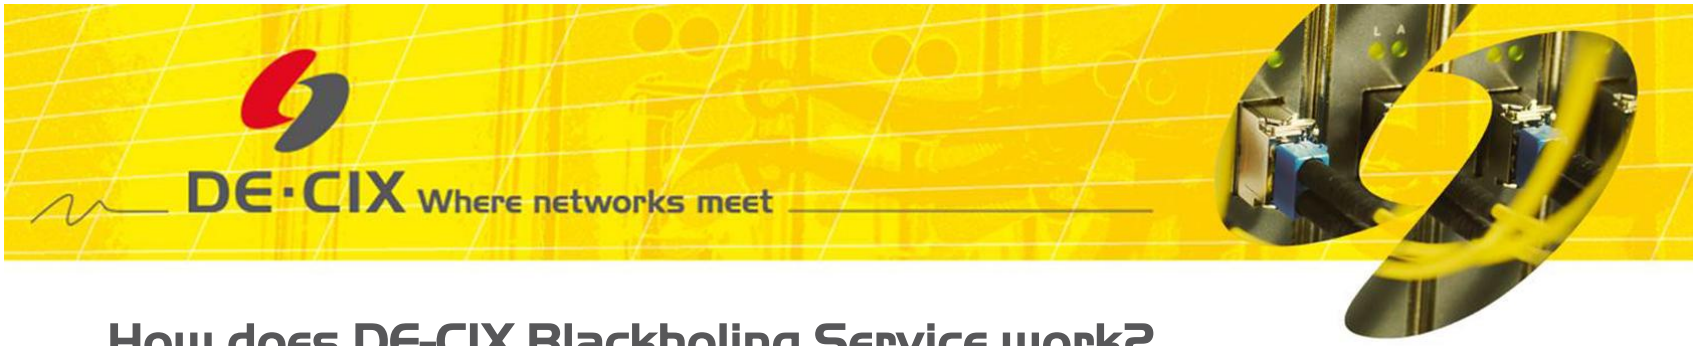

How does DE-CIX Blackholing Service work?

In standard conditions

- Customers advertise their prefixes with a next-hop IP address belonging to their AS, announced prefixes are from range of:
  - IPv4: /8 <= and <= /24
  - IPv6: /19 <= and <= /48

In case of attack

- Customers advertise their prefixes with a unique DE-CIX-provided *Blackhole Next-hop IP address* (BN)
  - IPv4: /8 <= **up to = /32** (if and only if the BN is set)
  - IPv6: /19 <= **up to = /128** (if and only if the BN is set)
- Further, same security checks apply as usual (whether the advertised prefix belongs to customer's ASN, etc.)

Current implementation of blackholing service

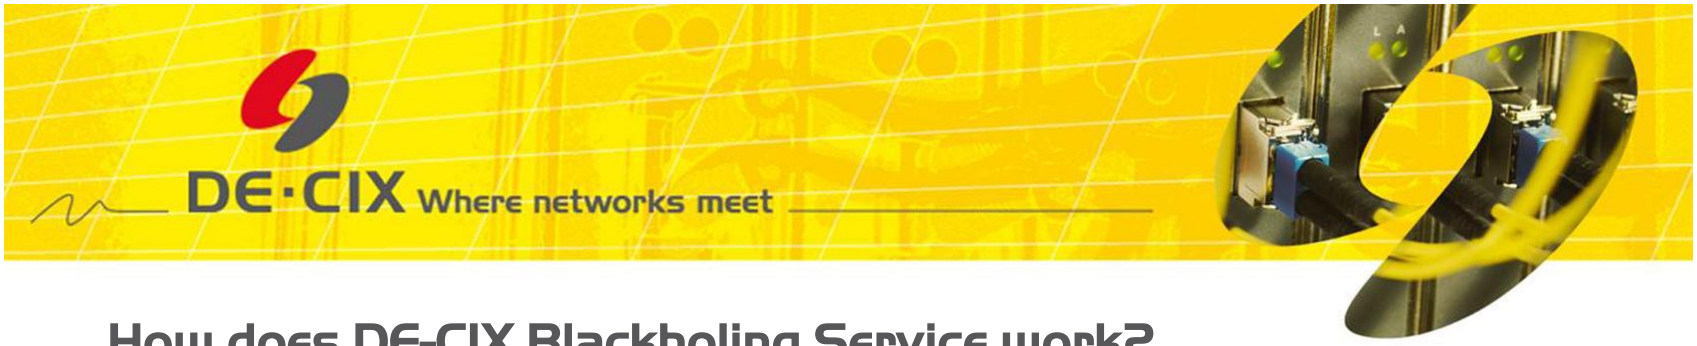

How does DE-CIX Blackholing Service work?

- L2 filtering
  - Blackhole Next-hop (BN) has a unique MAC address (determined by ARP for the BN IP address)
  - All frames with destination MAC address belonging to the BN are filtered ingress by the L2 ACL applied on all customer ports on our switches
- As a result, all traffic to the attacked and „blackholed“ prefix is discarded already on the switch, and hence victim‘s resources are protected

Current implementation of blackholing service

DE-CIX Blackholing Service – FAQs

- How many blackhole routes can I advertise?
  - Blackhole routes are included in the maximum number of advertised prefixes, hence number of your normal + blackhole routes should not exceed the allowed maximum
- How specific can the „blackholed“ prefix be?
  - The prefix can be as specific as /32 (IPv4) or /128 (IPv6)
- Do I have to pay for using the DE-CIX Blackholing Service?
  - No – use of blackholing is free of charge for customers
- At which locations is the DE-CIX Blackholing Service available?
  - The service is currently available only at DE-CIX Frankfurt

Opportunities associated with current implementation

- Current implementation has been designed to require little amount of changes at ISP side, if any, in order to be highly effective.
  - Let's assess this later on...
- Complements existing setups by closing an usually open configuration gap between peers.
- Traffic is blackholed at IXP, i.e. at a central point. This may open the road to more complex mitigation services.

Opportunities associated with current implementation

- Current implementation has been designed to require little amount of changes at ISP side, if any, in order to be highly effective.
  - Let's assess this later on...
- Complements existing setups by closing an usually open configuration gap between peers.
- Traffic is blackholed at IXP, i.e. at a central point. This may open the road to more complex mitigation services.

Opportunities associated with current implementation

- Current implementation has been designed to require little amount of changes at ISP side, if any, in order to be highly effective.
  - Let's assess this later on...
- Complements existing setups by closing an usually open configuration gap between peers.
- Traffic is blackholed at IXP, i.e. at a central point. This may open the road to more complex mitigation services.

Issues and pitfalls with current implementation

"Too large prefixes" discarded by peer

- **Assumption:** Peer follows best current practice(?) to filter incoming route advertisements based on prefix length, e.g.
  - IPv4: $/8 \leq p \leq /24$
  - IPv6: $/19 \leq p' \leq /48$
- **Issue:** Blackholing an IPv4 prefix $/24 \leq p \leq /32$ or an IPv6 prefix $/48 \leq p' \leq /128$ will not work.
- **Question:** Who follows this BCP?
- **Question:** Does this render the blackholing service useless?

Issues and pitfalls with current implementation

"Too large prefixes" discarded by peer

- **Assumption:** Peer follows best current practice(?) to filter incoming route advertisements based on prefix length, e.g.
  - IPv4: $/8 \leq p \leq /24$
  - IPv6: $/19 \leq p' \leq /48$
- **Issue:** Blackholing an IPv4 prefix $/24 \leq p \leq /32$ or an IPv6 prefix $/48 \leq p' \leq /128$ will not work.

- **Question:** Who follows this BCP?
- **Question:** Does this render the blackholing service useless?

Issues and pitfalls with current implementation

"Too large prefixes" discarded by peer

- **Assumption:** Peer follows best current practice(?) to filter incoming route advertisements based on prefix length, e.g.
  - IPv4: $/8 \leq p \leq /24$
  - IPv6: $/19 \leq p' \leq /48$
- **Issue:** Blackholing an IPv4 prefix $/24 \leq p \leq /32$ or an IPv6 prefix $/48 \leq p' \leq /128$ will not work.
- **Question:** Who follows this BCP?
- **Question:** Does this render the blackholing service useless?

Issues and pitfalls with current implementation

Receiver overwrites BGP next-hop attribtue

- **Assumption:** Peer follows best current practice(?) to overwrite BGP next-hop attribute of received prefix.
  - Cisco: `set ip next-hop peer-address`
  - Juniper: `set policy-options policy-statement peer-in term nexthop-peeraddr then next-hop peer-address`
- **Issue:** Blackholing service will not be effective, traffic will continue to be routed towards peer address.
- **Remark:** Only direct BGP sessions affected; sessions with DE-CIX route-servers *require* to not overwrite BGP next-hop attribute.
- **Question:** Who is affected by this? Who changed his config?
- **Question:** Do vendors support better/more-flexible next-hop filtering?

Issues and pitfalls with current implementation

Receiver overwrites BGP next-hop attribtue

- **Assumption:** Peer follows best current practice(?) to overwrite BGP next-hop attribute of received prefix.
  - Cisco: `set ip next-hop peer-address`
  - Juniper: `set policy-options policy-statement peer-in term nexthop-peeraddr then next-hop peer-address`
- **Issue:** Blackholing service will not be effective, traffic will continue to be routed towards peer address.

- **Remark:** Only direct BGP sessions affected; sessions with DE-CIX route-servers *require* to not overwrite BGP next-hop attribute.
- **Question:** Who is affected by this? Who changed his config?
- **Question:** Do vendors support better/more-flexible next-hop filtering?

Issues and pitfalls with current implementation

Receiver overwrites BGP next-hop attribtue

- **Assumption:** Peer follows best current practice(?) to overwrite BGP next-hop attribute of received prefix.
  - Cisco: `set ip next-hop peer-address`
  - Juniper: `set policy-options policy-statement peer-in term nexthop-peeraddr then next-hop peer-address`
- **Issue:** Blackholing service will not be effective, traffic will continue to be routed towards peer address.
- **Remark:** Only direct BGP sessions affected; sessions with DE-CIX route-servers *require* to not overwrite BGP next-hop attribute.
- **Question:** Who is affected by this? Who changed his config?
- **Question:** Do vendors support better/more-flexible next-hop filtering?

Issues and pitfalls with current implementation

Receiver overwrites BGP next-hop attribtue

- **Assumption:** Peer follows best current practice(?) to overwrite BGP next-hop attribute of received prefix.
  - Cisco: `set ip next-hop peer-address`
  - Juniper: `set policy-options policy-statement peer-in term nexthop-peeraddr then next-hop peer-address`
- **Issue:** Blackholing service will not be effective, traffic will continue to be routed towards peer address.

- **Remark:** Only direct BGP sessions affected; sessions with DE-CIX route-servers *require* to not overwrite BGP next-hop attribute.
- **Question:** Who is affected by this? Who changed his config?
- **Question:** Do vendors support better/more-flexible next-hop filtering?

Issues and pitfalls with current implementation

More specific blackholes

- **Assumption:** Peers accept blackhole prefixes up to /32 and /128 for IPv4 and IPv6, respectively.
- **Assumption:** Peers A and B share common downstream customer with prefix p under attack.
- **Assumption:** Peer B unconditionally announces blackhole prefix $p' > p$ (e.g. /32) and peer A does not.
- **Issue:** More specific wins. Peer A can no longer forward all packets to its customer.

Issues and pitfalls with current implementation

More specific blackholes

- **Assumption:** Peers accept blackhole prefixes up to /32 and /128 for IPv4 and IPv6, respectively.
- **Assumption:** Peers A and B share common downstream customer with prefix p under attack.
- **Assumption:** Peer B unconditionally announces blackhole prefix $p' > p$ (e.g. /32) and peer A does not.
- **Issue:** More specific wins. Peer A can no longer forward all packets to its customer.

Issues and pitfalls with current implementation

Unwanted RTBH in case of unauthorized route advertisement

- **Assumption:** Peer accepts spoofed next-hop attributes, i.e. prefixes with next-hop set to *BN*.
- **Assumption:** Prefix filtering at IXP does not scale; hence, trust in peers' announcements.
- **Assumption:** Peer *M* announces prefix *p* to peer *A* with next-hop set to *BN* unauthorized.
  - Advertisement of prefix *p* is unauthorized, iif *p* does not belong to set of prefixes induced by peer's AS-SET description.
- **Issue:** Unwanted and unauthorized RTBH of prefix *p*.
- **Remark:** Issue already existed before introduction of blackholing service.

Issues and pitfalls with current implementation

Unwanted RTBH in case of unauthorized route advertisement

- **Assumption:** Peer accepts spoofed next-hop attributes, i.e. prefixes with next-hop set to *BN*.
- **Assumption:** Prefix filtering at IXP does not scale; hence, trust in peers' announcements.
- **Assumption:** Peer *M* announces prefix *p* to peer *A* with next-hop set to *BN* unauthorized.
  - Advertisement of prefix *p* is unauthorized, iif *p* does not belong to set of prefixes induced by peer's AS-SET description.
- **Issue:** Unwanted and unauthorized RTBH of prefix *p*.

- **Remark:** Issue already existed before introduction of blackholing service.

Issues and pitfalls with current implementation

Unwanted RTBH in case of unauthorized route advertisement

- **Assumption:** Peer accepts spoofed next-hop attributes, i.e. prefixes with next-hop set to *BN*.
- **Assumption:** Prefix filtering at IXP does not scale; hence, trust in peers' announcements.
- **Assumption:** Peer *M* announces prefix *p* to peer *A* with next-hop set to *BN* unauthorized.
  - Advertisement of prefix *p* is unauthorized, iif *p* does not belong to set of prefixes induced by peer's AS-SET description.
- **Issue:** Unwanted and unauthorized RTBH of prefix *p*.

- **Remark:** Issue already existed before introduction of blackholing service.

Issues and pitfalls with current implementation

Corollary

- Blackholing service potentially conflicts with current BCPs
  - Prefix-length filters render blackholing service effectively useless if we assume single host blackholing the pre-dominant use case.
  - Overwriting BGP next-hop attribute to peer-address renders blackholing service useless.
- More-specific blackhole can affect service offered to customer.
- Unauthorized route advertisements can lead to remotely triggered blackhole.

$\Rightarrow$ Current implementation not feasible without change of peer configuration?

Issues and pitfalls with current implementation

Corollary

- Blackholing service potentially conflicts with current BCPs
  - Prefix-length filters render blackholing service effectively useless if we assume single host blackholing the pre-dominant use case.
  - Overwriting BGP next-hop attribute to peer-address renders blackholing service useless.
- More-specific blackhole can affect service offered to customer.
- Unauthorized route advertisements can lead to remotely triggered blackhole.

$\Rightarrow$ Current implementation not feasible without change of peer configuration?

Discussion

- How do we assess current service implementation?
- Can the service effectively be used without config change?
- Can we propose an alternative service implementation?
  1. Do we want to encourage next-hop spoofing at DE-CIX?
  2. Do we want BGP community support for blackhole prefixes?
  3. Do we want to set own internal blackhole next-hop and redistribute internally (i.e. drop traffic at ingress point)?
  4. Do we want more strict limitations on blackhole prefixes?
  5. Do we want increased monitoring and reporting of blackhole prefixes?
  6. How long can blackhole prefixes live?
- What about RFC 5575 *Dissemination of Flow Specification Rules*, aka. FlowSpec?
- What about OpenFlow / SDN?

Best practices for service utilization

- How do we think the service should be used by peers?
  1. Which prefixes should be announced?
  2. Which prefixes should be accepted?
- Should we draft a BCP document?

Effectiveness / impact

The following questions may sound academic, but ...

- Do we want to measure effectiveness of this service?
- How can we define *effectiveness*?
- How can we measure effectiveness of this servce?

Effectiveness / impact

The following questions may sound academic, but ...

- Do we want to measure effectiveness of this service?
- How can we define *effectiveness*?
- How can we measure effectiveness of this servce?

Alternative service implementation

- Do not encourage BGP next-hop spoofing.
- Add dedicated blackholing route-server (BHRS) and motivate peers to connect to it.
- Only announce blackhole prefixes via BHRS.
- Accept prefixes in the following ranges
  - IPv4: $/24 \leq p \leq /32$
  - IPv6: $/48 \leq p' \leq /128$
- Accept only **xx** prefixes per peer.
- Keep track of prefixes announced via BHRS
  - Add separate looking glass to BHRS.
  - Plot statistics on number of prefixes.
  - Automaticall generate mail to new mailinglist if new prefix is announced or announcement has stopped.

Examples of BGP peer configurations

Cisco IOS

...

Examples of BGP peer configurations

Juniper

...

Examples of BGP peer configurations

Brocade

...

CfP – Survey on attack detection and mitigation

Similar topics are content of publicly funded research projects.
We need some more insight, please participate in our survey!

- http://www.dasec.h-da.de/survey/netsec

da/sec
BIOMETRICS AND INTERNET-SECURITY
RESEARCH GROUP

eco

DE·CIX