WWW.WIDZEW.NET

# Implementation of **WANGUARD** software as a protection against DDOS attacks.

Piotr Okupski – STK TV-SAT 364
PLNOG 13 - 2014

# TV /ISP Cable Company

TV-SAT 364

- Over 7000 customers
- Services :TV, Internet, Phone
- Network Equipment: Juniper, Cisco , D-Link
- Over 20 years on market


# Piotr Okupski

- CTO
- Administration of network and servers
- Planning and development
- pl.**linkedin**.com/in/**piotrokupski**

# Agenda

- Case study
- **WANGUARD deployment scenarios**
- Hardware requirements and hardware available on the market
- Server and Linux tuning
- Configuration of **WANGUARD and Juniper MX**
- How fast is **WANGUARD?**
- Additional Juniper Routing Engine protection
- Whats new of DDoS front?

# Case study– Atack!

1. DDoS Attack on our network:
   - Mainly UDP Flood (NTP,DNS)
   - 3-4 Times a day, for 3-5 minutes for 2 weeks

2. Prepairing to defend
3. System implementation
4. BGP tests and blackholing
5. System tuning

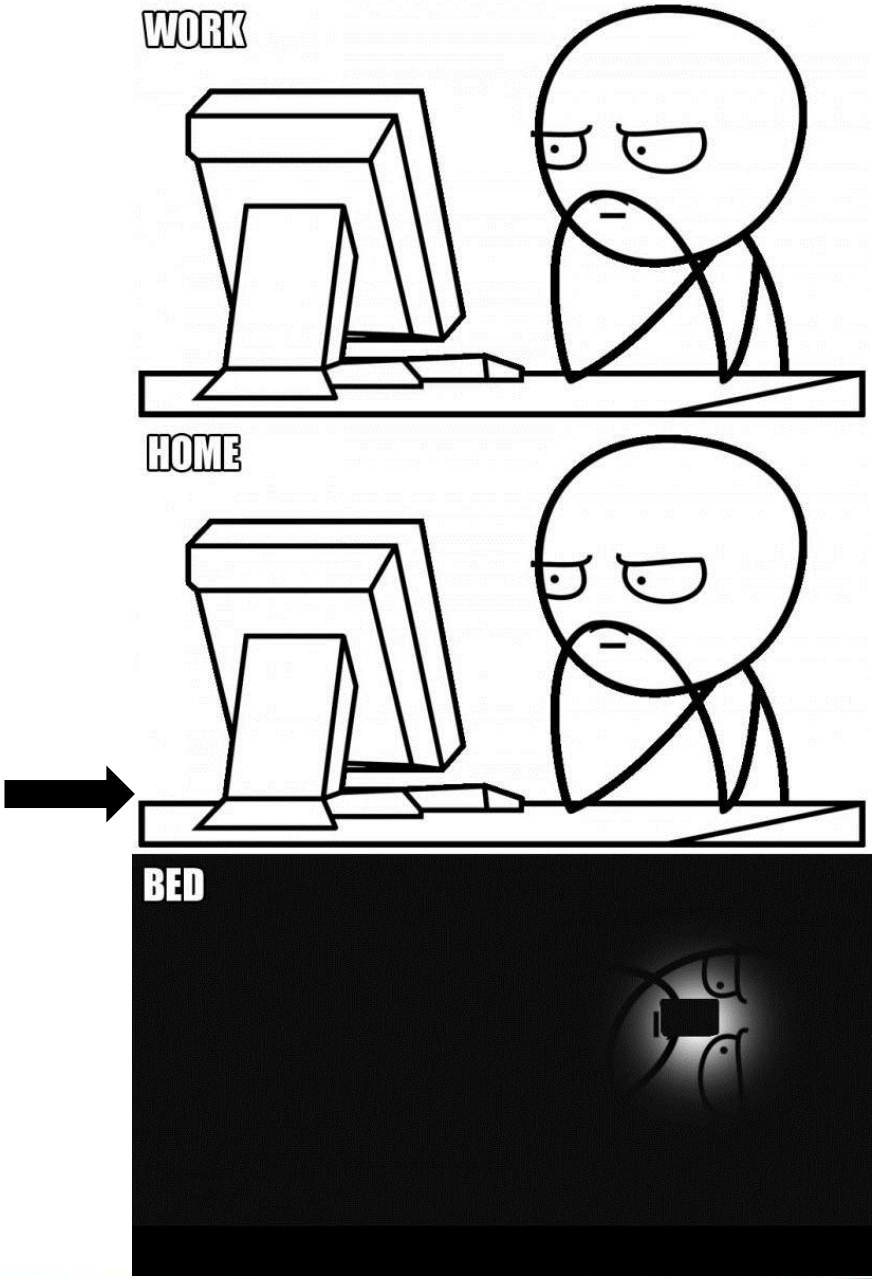

# Our attacks

**Anomaly #90**

**Feb 16 15:23:47 - 15:27:07 (3m 20s)**

| Anomaly | Value | Affected | Traffic | Severity |
|---|---|---|---|---|
| Incoming TOTAL pkts/s > 20.0k | Highest: 82.6k<br>Last: 195.9k | External Zone ( ) | 82.6k pkts/s<br>873.7 Mbits/s | |

**Details**

| Sensor | Response | IP Zone | Threshold | Totals | Actions |
|---|---|---|---|---|---|
| WAN | BGP Blackhole (Blackhole-Wanguard, Raport) | WAN-Classes ( | 20.0k | 8.9 Mpkts 93.4 Gbits | |

82.6k pkt/s

873.7 Mbit/s

**Anomaly #129**

**Jun 25 18:28:47 - 18:28:57 (10s)**

| Anomaly | Value | Affected | Traffic | Severity |
|---|---|---|---|---|
| Incoming TOTAL pkts/s > 20.0k | Highest: 126.7k<br>Last: 9.4M | External Zone ( ) | 126.7k pkts/s<br>454.8 Mbits/s | |

**Details**

| Sensor | Response | IP Zone | Threshold | Totals | Actions |
|---|---|---|---|---|---|
| WAN | BGP Blackhole (Blackhole-Wanguard, Raport) | WAN-Classes ( | 20.0k | 1.2 Mpkts 4.5 Gbits | |

126.7k pkt/s

454.8 Mbit/s

**Anomaly #109**

**Mar 31 20:25:27 - 20:25:32 (5s)**

| Anomaly | Value | Affected | Traffic | Severity |
|---|---|---|---|---|
| Incoming TOTAL pkts/s > 20.0k | Highest: 159.6k<br>Last: 352.5k | External Zone ( ) | 159.6k pkts/s<br>594.1 Mbits/s | |

**Details**

| Sensor | Response | IP Zone | Threshold | Totals | Actions |
|---|---|---|---|---|---|
| WAN | BGP Blackhole (Blackhole-Wanguard, Raport) | WAN-Classes ( | 20.0k | 1.2 Mpkts 4.4 Gbits | |

159.6k pkt/s

594.1 Mbit/s

# Wanguard (Annual subscriptions)

| WAN SIGHT | WAN GUARD Sensor | WAN GUARD Filtr |
|---|---|---|
| WAN **Sight Sensor** | WAN Guard **Sensor** | WAN Guard Filter |
| • Sensor flow/pcap<br>• Traffic graphs, stats<br>• **Doesn't detect attacks**<br>• Realtime traffic analysis | • Sensor flow/pcap<br>• Traffic graphs, stats<br>• Realtime traffic analysis<br>• **Traffic anomalies detection** | • **Side filtering of malicious traffic** |
| 345 $ | 595 $ | 995 $ |

# Wanguard deployment scenarios

## Sensor and Filter with BGP Routing

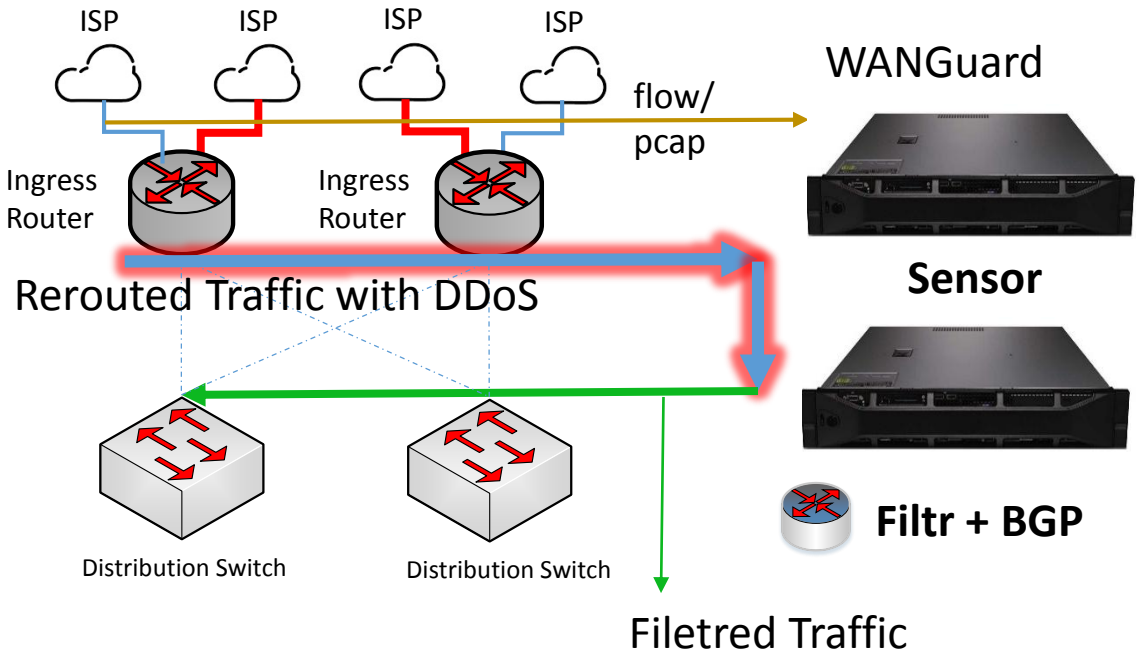


### Scheme plan:

- After attack detection Wanguard send BGP update to router to reroute traffic to filter.
- Filter is dropping all traffic with attack signatures and reroutes back only valid traffic.

# Wanguard deployment scenarios

## Sensor and Filter with BGP router (All in one)

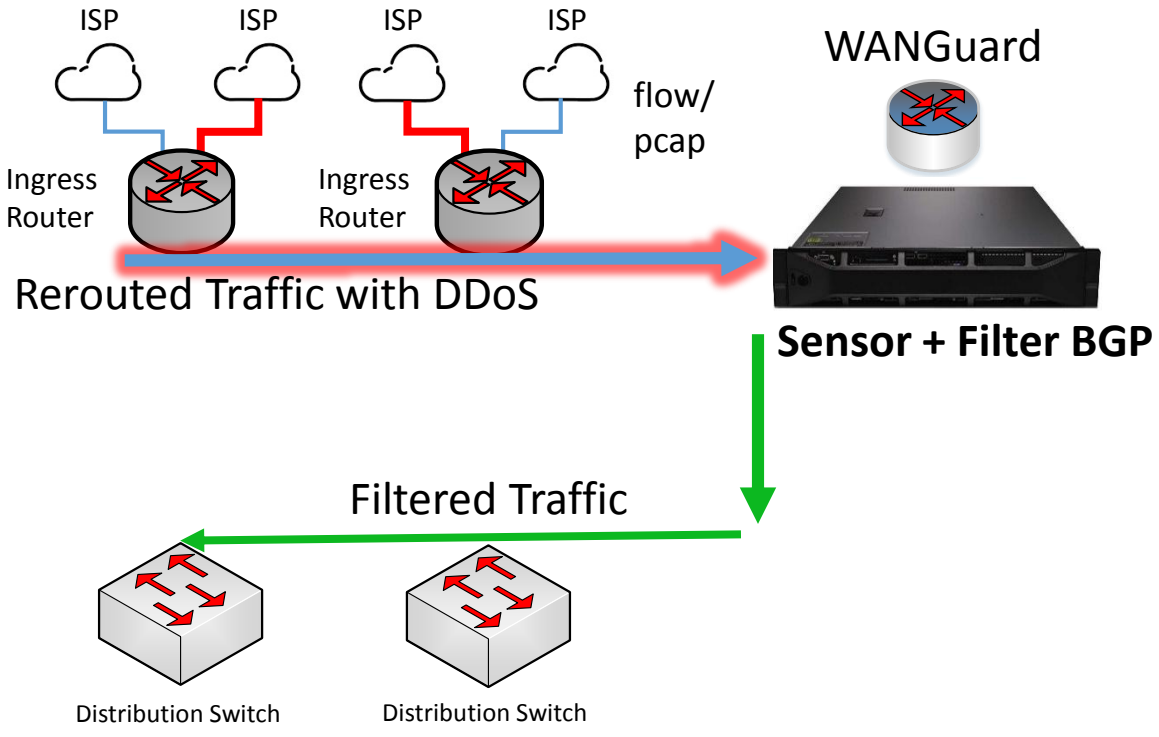


## Deploment scheme

- Sensor and filter and Quagga router are installed on the same server simplifying routing process and saving one server.

# Wanguard deployment scenarios

## Sensor and Filter as L2 bridge

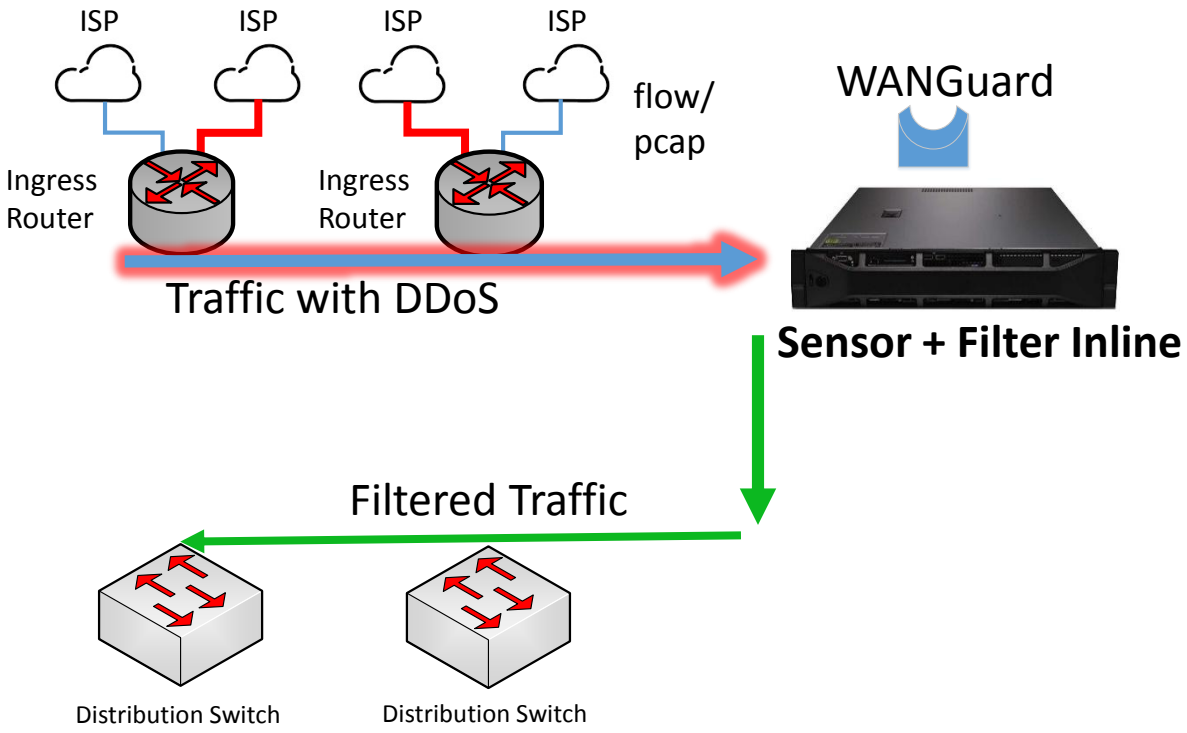


## Deployment scheme

- Sensor and filter are working INLINE filtering all the traffic that's passing through the server.

# Wanguard deployment scenarios

## Sensor as passive sniffer with BGP for RTBH

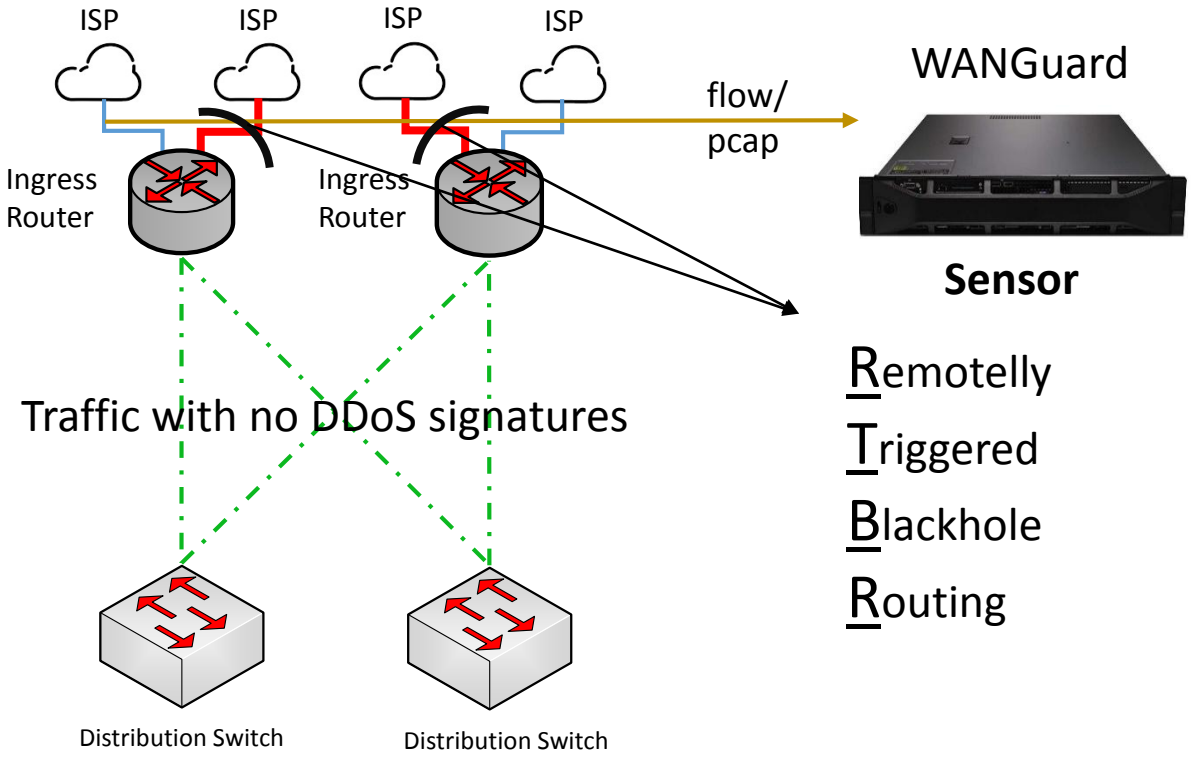


Remotelly
Triggered
Blackhole
Routing

### Deployment sheme:

- Server is gathering copy of all traffic via RSPAN/mirror port and avoiding Netflow delay.
- Sensor is sending BGP update with prefix to our routers that's going to be blackholed by our upstream providers.

# RTBH How does it work?

ATTACK

WANGUARD SERVER
(Sniffer & Quagga[BGP])

BGP Prefix 1.1.1.1/32
[42830:666]

SNIFFED TRAFFIC

ATMAN

BGP Prefix 1.1.1.1/32
[3257:2666]

Internet

TV-SAT 364

BGP Prefix 1.1.1.1/32
[24724:666,3257:2666]

BGP Prefix 1.1.1.1/32
[24724:666]

GTT

# Hardware requirements for Wanguard

| Appliance | Do 1 Gb/s | Do 10 Gb/s |
|---|---|---|
| **Deployment scheme** | In-line lub out-of-line | Recomended out-of-line |
| **CPU:** | 2.5 GHz dual-core Xeon | 2.8 GHz quad-core Xeon |
| **RAM:** | 2 GB | 8 GB |
| **Network:** | 2 x Gigabit Ethernet | 1 x 10 GbE card (Intel 82599/x520/x540)<br>1 x Gigabit Ethernet<br>lub Karty: Napatech / Emulex |
| **OS:** | RHEL / CentOS 5, RHEL / CentOS 6, OpenSUSE 12, Debian 6 / 7, Ubuntu Server 12 | RHEL / CentOS 5, RHEL / CentOS 6, OpenSUSE 12, Debian 6 / 7, Ubuntu Server 12 |
| **HDD:** | 10 GB (including system)<br>+ extra GB for flows and grahs | 10 GB (including system)<br>+ extra GB for flows and grahs |

# Network cards 1-100 Gb/s

| Company | Card | Symbol | Price* | Drivers options (ex. 0 % CPU) |
|---|---|---|---|---|
| Emulex (Endance) (1-10 GbE) | 2 port 1GbE | DAG 7.5G2 | 800 -1200 $ | • Manufacturer<br>• PF_RING |
| | 2 port 10 GbE | DAG 9.2X2 | 14000 $ | |
| **Intel (1-10)** | 2 port 10GbE | x520/x540 | 400-600 $ | Alt.PF_RING – 250€/MAC |
| **Mellanox** | 2 port 10GbE | MCX312A-XCBT | 400$ | |
| | 2 port 40 GbE | MCX314A-BCBT | 850$ | |
| **Myricom (10GbE)** | 1 port 10GbE | 10G-PCIE-8B-S | 400 $ | Sniffer10Gv2/v3 – 260$/530$** |
| | 2 port 10GbE | …-8B2-2S | 800 $ | |
| **Napatech (1-100 GbE)** | 4-port 1GbE | NT4E2-4-PTP | 5274 $* | • Manufacturer<br>• PF_RING |
| | 2-port 10GbE | NT20E2-PTP | 13,562 $* | |
| | 1-port 40GbE | NT40E2-1 | 13,562 $* | |
| | 1 port 100GbE | NT100E3-1-PTP | 33,908 $* | |

* *MSRP Price* 08.2014 (update)

** Wanguard 5.5

# System Tuning – BIOS

| Option | | | Value |
|---|---|---|---|
| **General** | Operating Mode/PWR Profile | | Maximum Performance |
| **Procesor** | | C-States | Disabled |
| | | Turbo mode | Enabled/Disabled (zależy od modelu) |
| | | C1E | Disabled |
| | | HT | Disabled |
| | | Intel VT-x/i/c/d | Disabled |
| **Memory** | | Memory Pre-Failure Notification | Disabled |
| Memory Speed | | | Maximum Performance |
| Memory Channel mode | | | Independent |
| Node Interleaving | | | Disabled/NUMA |
| Thermal Mode | | | Performance |
| **Power** Profile | | | Maximum Performance |

Tuning for low latency systems (high performance) np.: DELL/HP/IBM [ http://goo.gl/cxIvL1 ]

# System tuning – Linux part. 1

1. **IRQ Balancer** should be turned off and removed from system (**1**)
2. RX/TX queues must be pinned for each CPU (**2**)
   ( **1** i **2** only if we are NOT! using PF_RING and other divers )
3. **CPU Affinity should be set for each WANSniff** proces
   (by hand or script)
4. **CPU Scalling Governor** set as ondemand [or performance]

```
#cat /sys/devices/system/cpu/cpu0/cpufreq/scaling_governor
ondemand
```

**All tuning must be adequate to your network need (packets/volume etc)**

Source: www.mellanox.com/related-docs/prod_software/Performance_Tuning_Guide_for_Mellanox_Network_Adapters.pdf

# System tuning – Linux part. 2

5. **I/O Scheduler** is set for deadline for SSD drives at grub level

```
title Red Hat Enterprise Linux Server (2.6.18-8.el5)
        root (hd0,0)
        kernel /vmlinuz-2.6.18 ro root=/dev/sda2 elevator=deadline
```

```
# cat /sys/block/sda/queue/scheduler
noop [deadline] cfq
```

6. Export for graphs is set for ramdrive (tmpfs) or to SSD drive

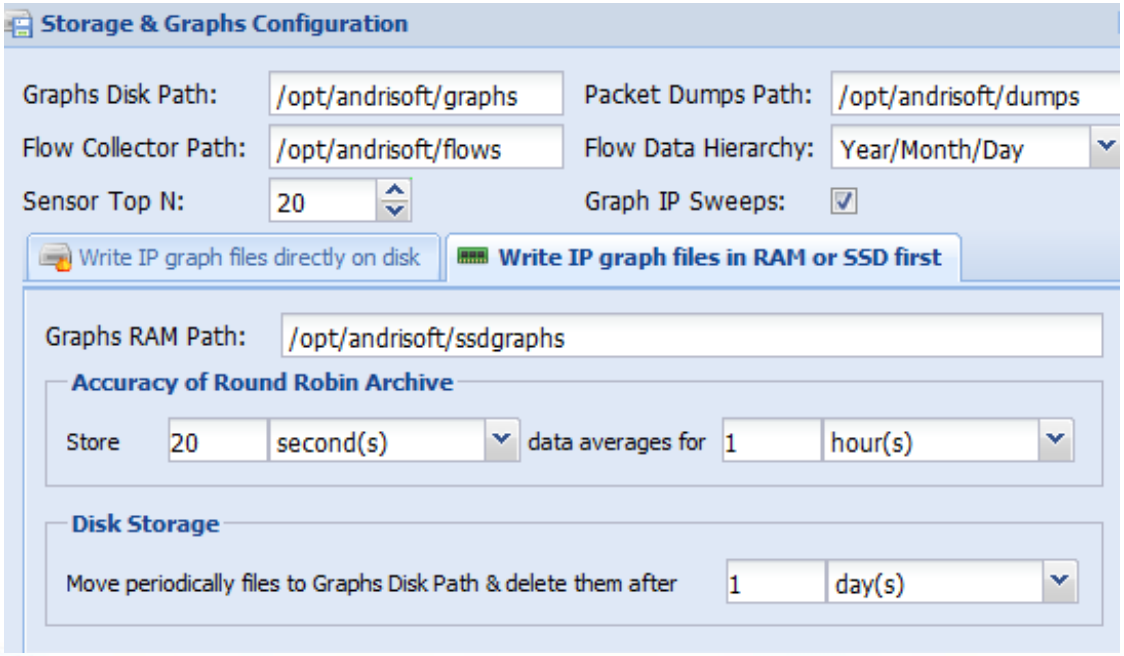

# Script to pin **WANsniff** per CPU Core

**Wanguard 5.4**
(no built-in CPU Affinity mechanism)

```
#!/bin/bash
CPU_WG=2
WG_SNIFF=`pgrep WANsniff`
for wpid in $ WG_SNIFF
do
 taskset -cp $CPU_WG $wpid
 CPU_WG=$[$CPU_WG +1]
done
```

**Wanguard 5.5**
( version available in 2-3 months time)

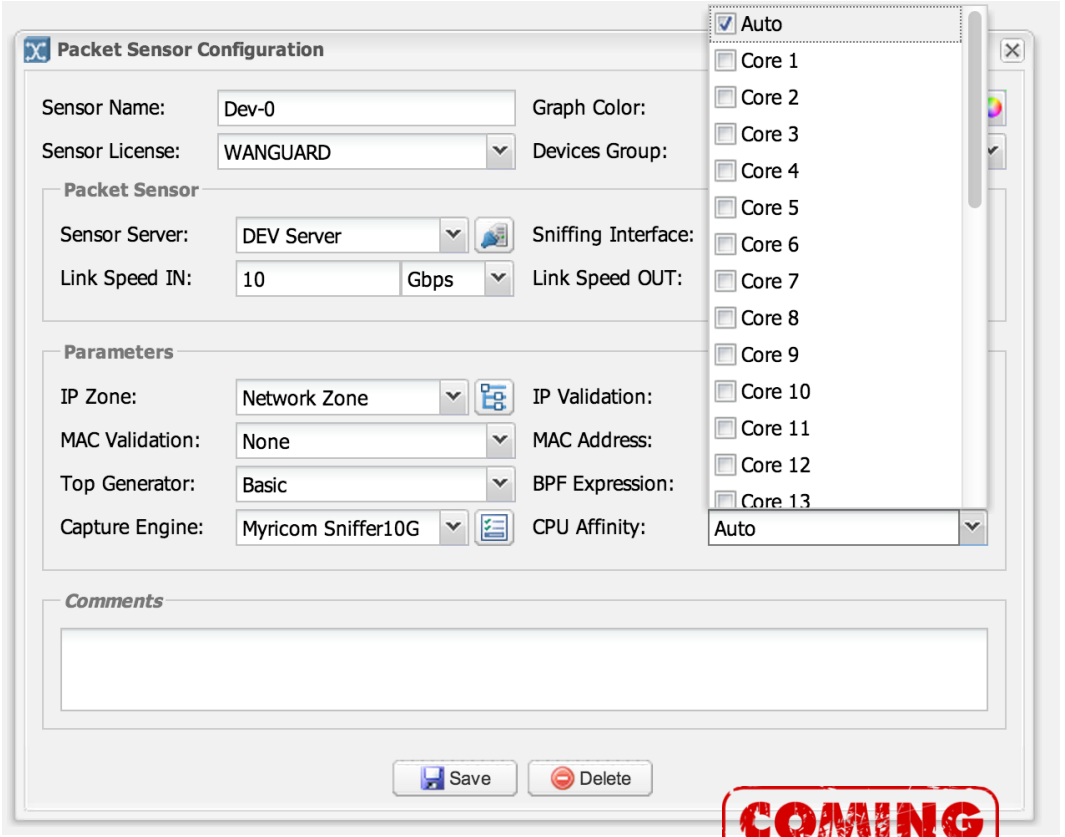

# Configuration scheme

**WANGuard Server**
Sniffer + Quagga BGP

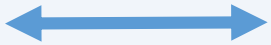

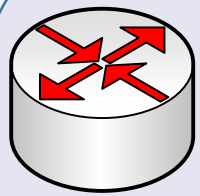

TV-SAT 364

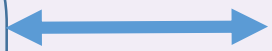

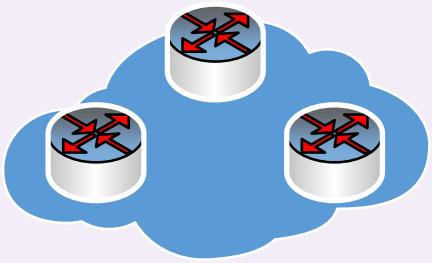

ISPs and IX-es

**Quagga Configuration– Juniper 1/2**
[bgpd.conf] 1

**Quagga Configuration– Juniper 2/2**
[bgp group IBGP-Wanguard] 2
[policy-statement import-Wanguard]

**Juniper Policy 1/2** 3
[protocols bgp group EBGP-ATMAN]

**Juniper Policy 2/2** 4
[policy-statement export-myAS-with-blackholing ]

# Quagga configuration – Juniper 1/2

1

## Konfiguracja QUAGGA

```
router bgp 42830
bgp router-id 10.0.0.2
neighbor 10.0.0.1 remote-as 42830
neighbor 10.0.0.1 description R1
neighbor 10.0.0.1 next-hop-self
neighbor 10.0.0.1 soft-reconfiguration inbound
neighbor 10.0.0.1 distribute-list nothing-in in
neighbor 10.0.0.1 route-map WANGUARD-Filter-out out
!
access-list nothing-in deny any
route-map WANGUARD-Filter-out permit 10
set community 42830:666
```

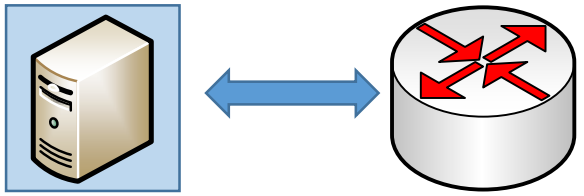


- Quagga - same AS as on our router.
- iBGP session is created.
- Next-hop-self is added to fix problem when importing prefix from iBGP to eBGP.

**666** – Blackhole Community

# Quagga configuration– Juniper 2/2

2

## Juniper

```
[protocols bgp group IBGP-Wanguard]
type internal;
description Wanguard-Blackholing;
family inet {
    unicast {
        rib-group only-inet.0;
    }
}
peer-as 42830;
neighbor 10.0.0.2 {
    no-advertise-peer-as;
    import import-Wanguard;
    export no-export;
}
```

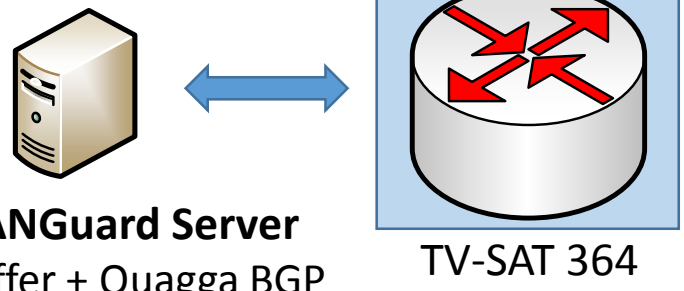


- Nie eksportujemy nic do Wanguard-a.
- Używamy polityki **import-Wanguard**, aby sterować co będziemy dalej robić z wykrytymi prefiksami podczas ataku.

# Juniper Policy 1/2

3

Juniper

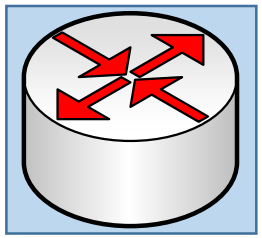

TV-SAT 364

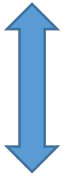

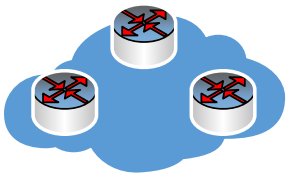

ISPs i IX-es

```
[policy-options policy-statement import-Wanguard]
term local_exceptions {
   from {
      protocol bgp;
      prefix-list local-exclude-from-blackholing;
            }
   then reject; }
term only_32_prefixes {
   from {
      protocol bgp;
      community com-wanguard;
      route-filter 2.2.2.0/22 prefix-length-range /32-/32;
      route-filter 3.3.3.0/20 prefix-length-range /32-/32;
   }
   then {
      community delete com-wanguard;
      community add blackhole-ATM;
      community add blackhole-TINET;
   accept;   }
            }
term last-deny-all {  then reject; }
```

# Juniper Policy 2/2



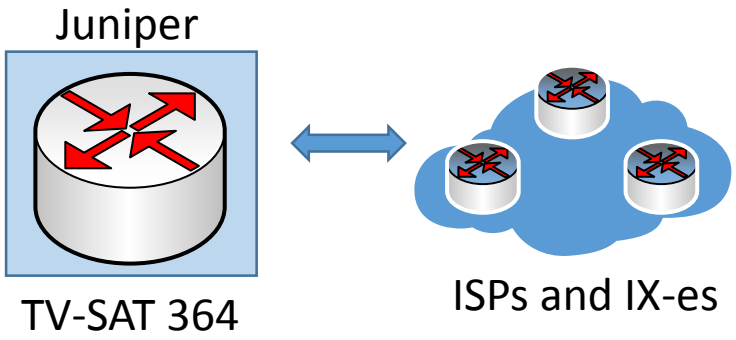


## Juniper BGP

**[edit protocols bgp group EBGP-ATMAN ]**
type external;
description ATMAN-World;
import import-atm-glob;
family inet {
  unicast {
    rib-group only-inet.0;
  }
}
export
**export-myAS-with-blackholing;**
peer-as 24724;
neighbor 193.x.x.x;

## Juniper BGP

**[edit policy-options policy-statement export-myAS-with-blackholing ]**
term blackhole {
from {
protocol bgp;
community [ **blackhole-ATM blackhole-TINET** ];
  }
  then accept;
}
term my-PI-prefixes {
  from {
    prefix-list PI-SPACE;
  }
  then accept;
}
term no-transit {
  then reject;
}
**[edit policy-options]**
community **blackhole-ATM** members 24724:666;
community **blackhole-TINET** members 3257:2666;
community **com-wanguard** members **42830:666**
community blackhole-providers members
[ target:24724:666 target:3257:2666];

# How fast is WANGUARD ?

WANGUARD CONSOLE

Anomalies Dashboard | BGP Prefixes | Console | Anomalies | Packet Analyzers | All Sensors - Extended Tops

Actions: Clear Active Announcements

Rows Filtering Expression

| № | Anomaly# | BGP Connection | IP Address | Mask | From | Until |
|---|---|---|---|---|---|---|
| 145 | - | BGP-Wanguard | 9.9.9.9 | 32 | 2014-08-31 15:26:14 | 2014-08-31 15:27:14 |

15:26:14

| | | | | | |
|---|---|---|---|---|---|
| 570375 | Console | BGP Interface | DEBUG | Sending BGP announcement # 145 requested by "admin" | 2014-08-31 15:26:15 |

15:26:15

**[Start of the update to Juniper]**

Aug 31 **15:26:18**.313185 BGP RECV 10.0.0.52+41386 -> 10.0.0.1+179

15:26:18

**[Update to peers]**

Aug 31 15:26:18.315247 bgp_send: sending 59 bytes to 7x.x.x.x (External AS 3257)

Aug 31 15:26:18.315484 BGP SEND 9.9.9.9/32

Aug 31 15:26:18.316102 bgp_send: sending 59 bytes to 1x.x.x.x (External AS 24724)

Aug 31 **15:26:18**.316331 BGP SEND 9.9.9.9/32

**[Last action:15:26:18.316331]**

00:00:00.003146 = **3.146 ms – BGP, Total time**

**4 seconds !**

# Additional security Juniper RE (1)

**Day One Book: Securing The Routing Engine on M/MX/T Series**
– Douglas Hanks Jr. [ http://goo.gl/648hrU ]

## Juniper Routing Engine (lo0)

```
lo0 {
    unit 0 {
        family inet {
            filter {
             input-list [ accept-bgp accept-common-services accept-established discard-all ];
            }
        }
        family inet6 {
            filter {
             input-list [ accept-v6-bgp accept-v6-common-services accept-established-v6 discard-
v6-all ];
            }
        }
    }
}
```

# Additional security Juniper RE (2)

## Juniper accept-bgp filter

```
show firewall filter accept-bgp {
apply-flags omit;
term accept-bgp {
from {
 source-prefix-list {
            bgp-neighbors;
            bgp-neighbors-logical-systems;
}
destination-prefix-list {
            router-ipv4;
            router-ipv4-logical-systems;
}
            protocol tcp;
            port bgp;
}
then {
            count accept-bgp;
            accept;
}
}
```

# Additional security Juniper RE(3)

## Juniper prefix lists

**show policy-options prefix-list router-ipv4**
apply-path “interfaces <*> unit <*> family inet address <*>”;

**show policy-options prefix-list bgp-neighbors**
apply-path “protocols bgp group <*> neighbor <*>”;

**show policy-options prefix-list bgp-neighbors-logical-systems**
apply-path “logical-systems <*> protocols bgp group <*> neighbor <*>”;

**show policy-options prefix-list router-ipv4-logical-systms**
apply-path “logical-systems <*> interfaces <*> unit <*> family inet address <*>”;

**Option for IPv6**

**show policy-options prefix-list router-ipv6**

apply-path “interfaces <*> unit <*> family inet address **<*:*>**”;

# Additional security Juniper RE (4)

## Juniper apply-path in action

**show configuration policy-options prefix-list router-ipv4 | display inheritance**
##
## apply-path was expanded to:
##  x.x.x.x/30;
##  x.x.x.x/22;
##  x.x.x.x/30;
##  x.x.x.x/30;
##  x.x.x.x/23;

## Juniper BGP hold-time (time to tear down BGP session)

**[edit protocols bgp group EBGP-example]**
type external;
description EBGP_Session_Example;
hold-time 90; [ Default 90 seconds, 0 disables keepalive]
import import-policy;

# What's new of DDoS front?

**FlowSpec RFC 5575**

1. Juniper - Junos 15.+ COMING SOON!
   suport for ISSU/NSR, Redirect i IPv6 – in draft
2. eXaBGP – FlowSpec READY!
3. WANGUARD FlowSpec suport – COMING SOON!
   only scripts + API with WANGUARD Filter.
4. Firewall on Demand –
http://code.grnet.gr/projects/flowspy/

# Polish IX-es and RTBH support

| IX Name | RTBH ? | Webpage |
|---|---|---|
| EPIX | YES √ | www.epix.net.pl |
| KIX | YES √ | tanielacze.pl |
| PLIX | YES √ | www.plix.pl |
| THINX | YES √ | www.thinx.pl |
| TPIX | YES √ | www.tpix.pl |

**Piotr Okupski**

okupski at widzew.net