

RIPE NCC Days Sofia, 28 June 2023



RIPE NCC

RIPE NETWORK COORDINATION CENTRE



QoS and DDoS mitigation in IXP

Quality of Service

Different traffic treatment for different services

- **Best-effort** - default
- **Assured forwarding** - reserved bandwidth
- **Expedited forwarding** - low latency (priority)
- **Network Control** - priority and reserved bandwidth

Input **classification** and output **queueing**

What it has to do in neutral IXP infrastructure



QoS on Peering only Member port

Output Queueing

Best effort for everything

Network Control for BGP (and ARP)



QoS on Peering + Multicast

Multicast is used for real-time Audio/Video transport
Highly sensitive for packet loss and variable delay (jitter)

Output Queueing

Best effort for Peering

Network Control for BGP

Expedited forwarding for Multicast



QoS on Multi-service Member port

Private VLANs normally carry more important traffic than peering

Output Queueing

Best effort for Peering

Network Control for BGP

Expedited forwarding for Multicast

Assured forwarding for Private VLANs



QoS with DDoS mitigation

Reserve zero bandwidth for possible DDoS traffic

Output Queueing

Best effort for possible DDoS

Assured forwarding for other Peering

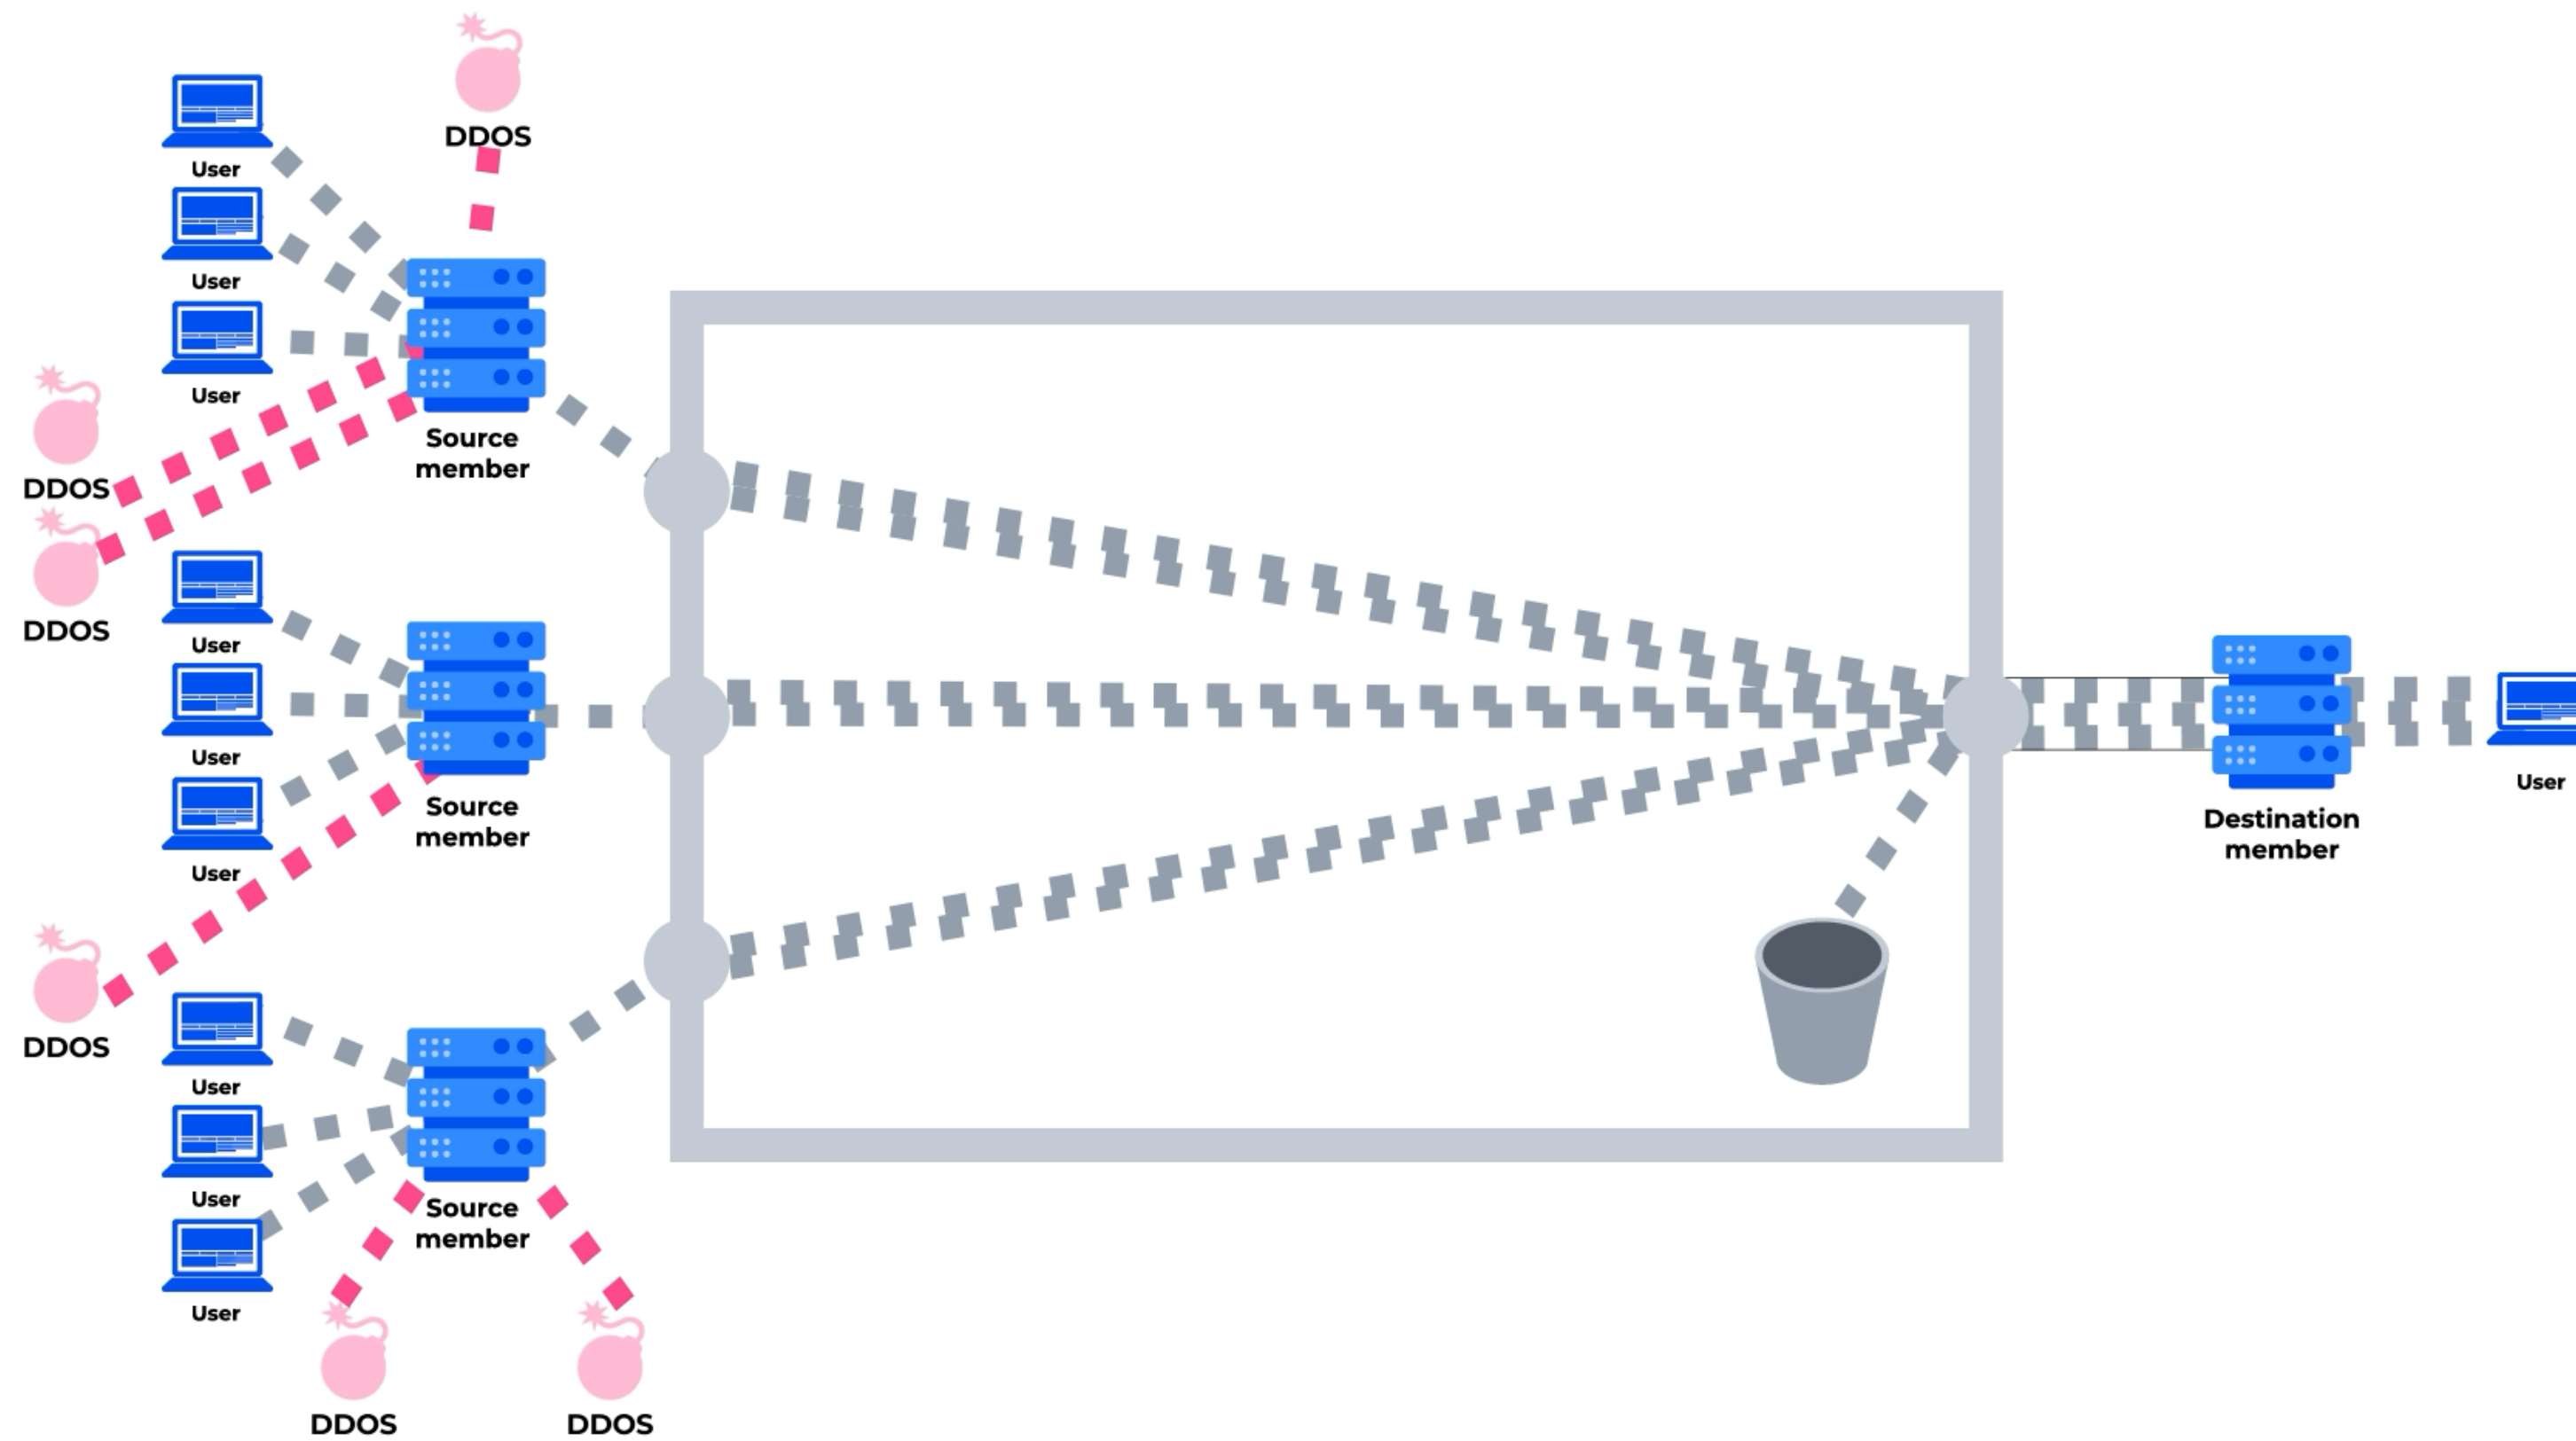
Network Control for BGP

Expedited forwarding for Multicast

Assured forwarding for Private VLANs



Congestions on Member ports



- Routing change
 - policy
 - public peering as backup of private peering
- Increased traffic
 - live events
 - software updates
- DDoS attacks



Detecting Congestion Cause

SFlow / Port Mirroring

PMacct

ElasticSearch

Zabbix



DDoS Attacks

Low-and-Slow DDoS attack

protocol and application layer

High packet rate DDoS attacks

TCP syn, DNS, HTTP/S
spoofed traffic

Solutions: hardware/software solutions, BCP.38

Volumetric DDoS attack

overflow transport links to the victim

Amplification
Bot networks



DDoS Patterns

Application	Protocol	Port
Invalid	UDP	0
Chargen	UDP	19
DNS	UDP	53
NTP	UDP	123
SNMP	UDP	161
U Discovery	UDP	10001
Memcache	UDP	11211
SSDP	UDP	19000

Amplification

- Request from spoofed IP address of the victim
- Large response to target

Most attacks last less than 2 minutes
Multiple protocols



DNS DDoS attacks

DNS is important but low bandwidth traffic
(Less than 10 Mbps on whole IXP)
Multiple Gbps during attack

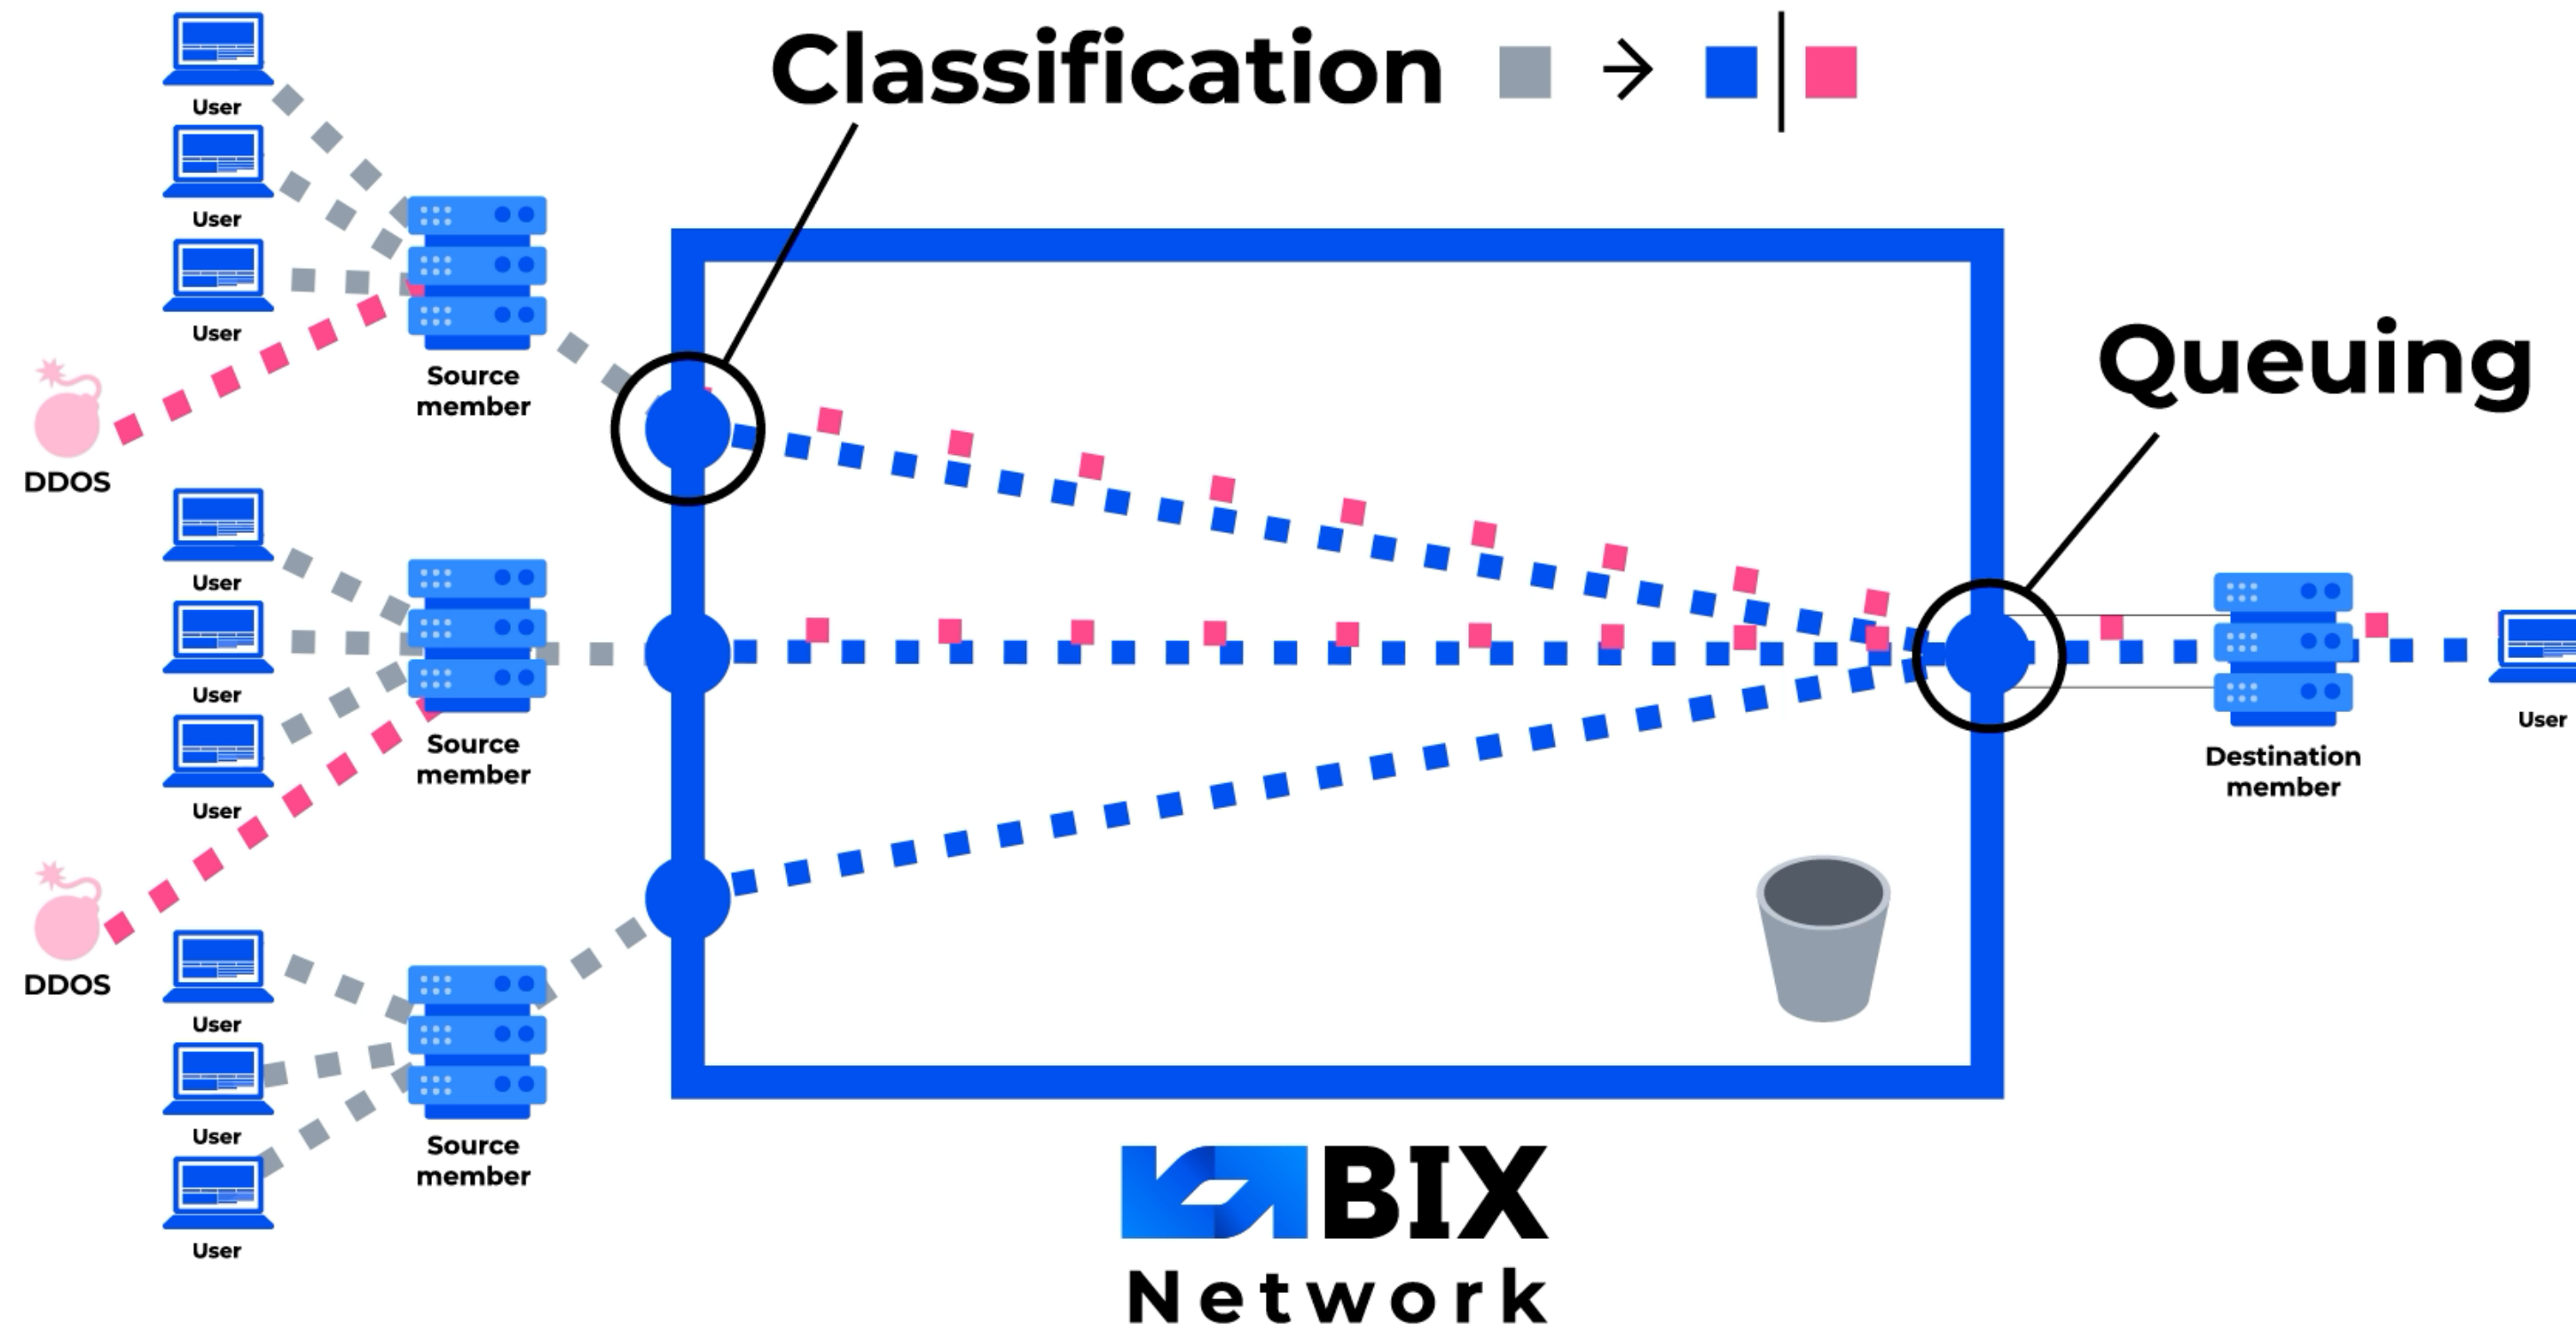
Solution: Policing

Classify UDP with source 53 exceeding
50 Mbps on all ports as potential DDoS

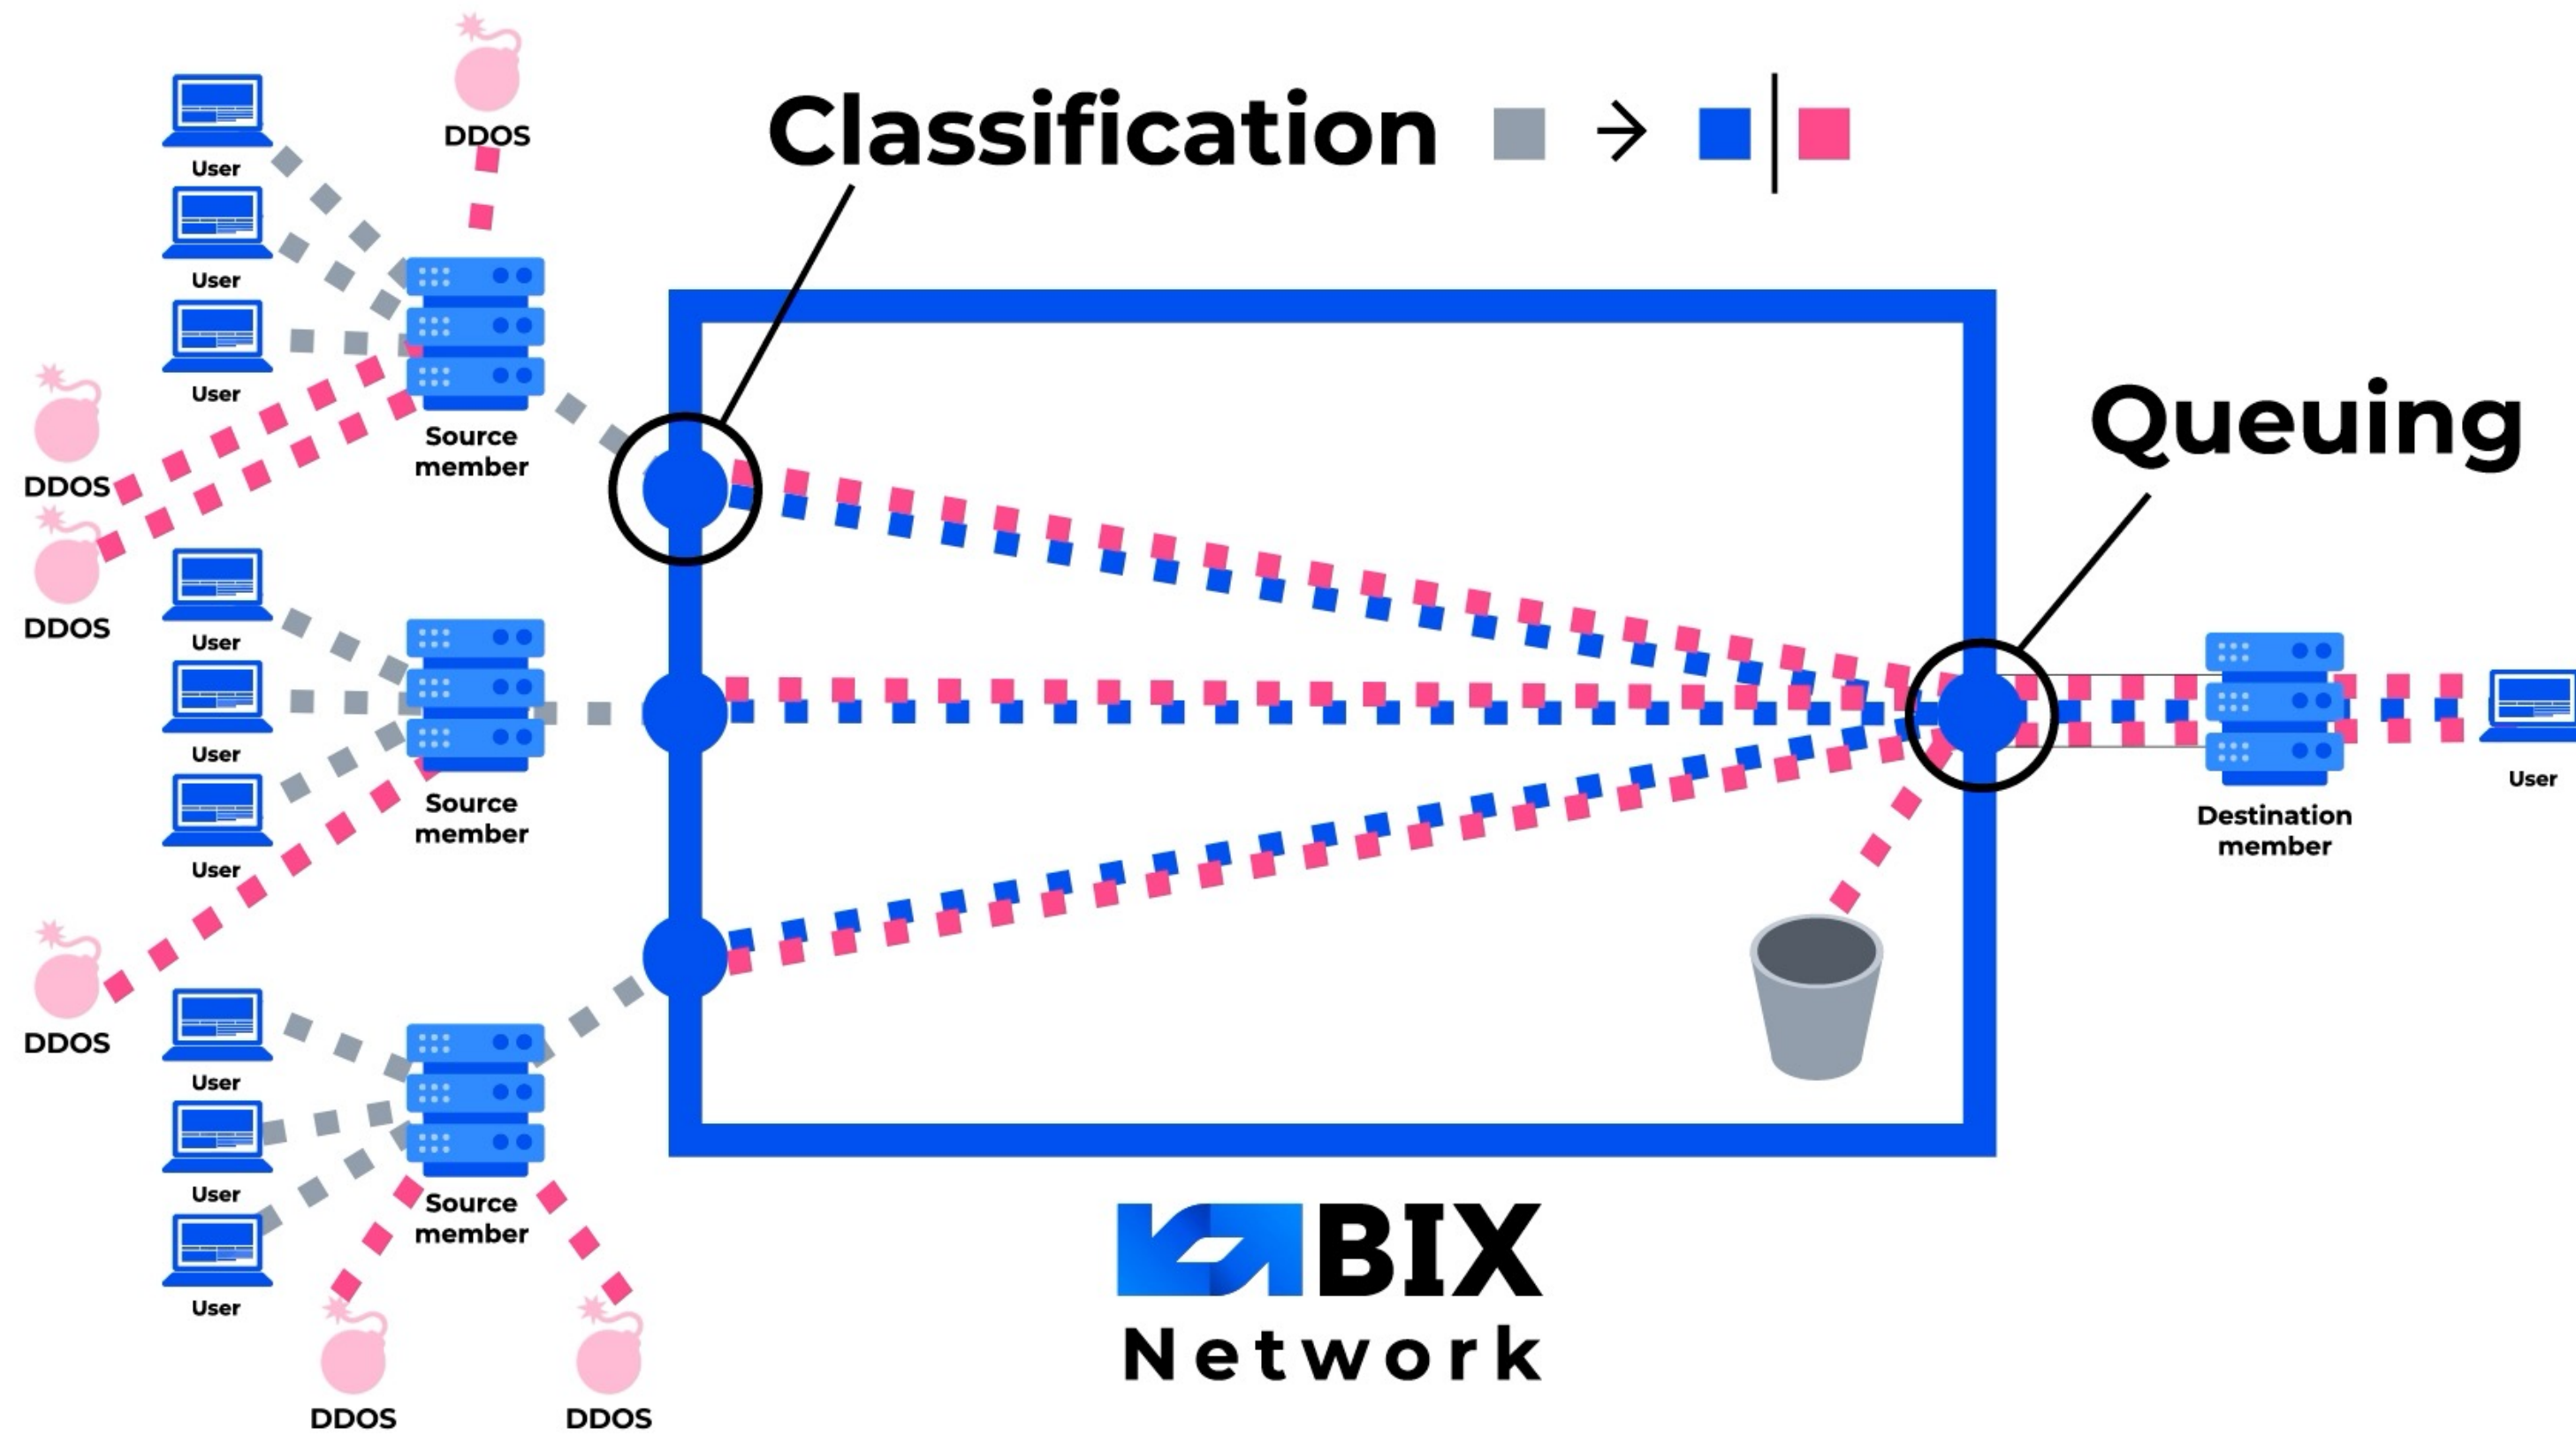
Same solution for other patterns




DDoS attack w/o congestion




DDoS attack causing congestion



Real Attack in Monitoring



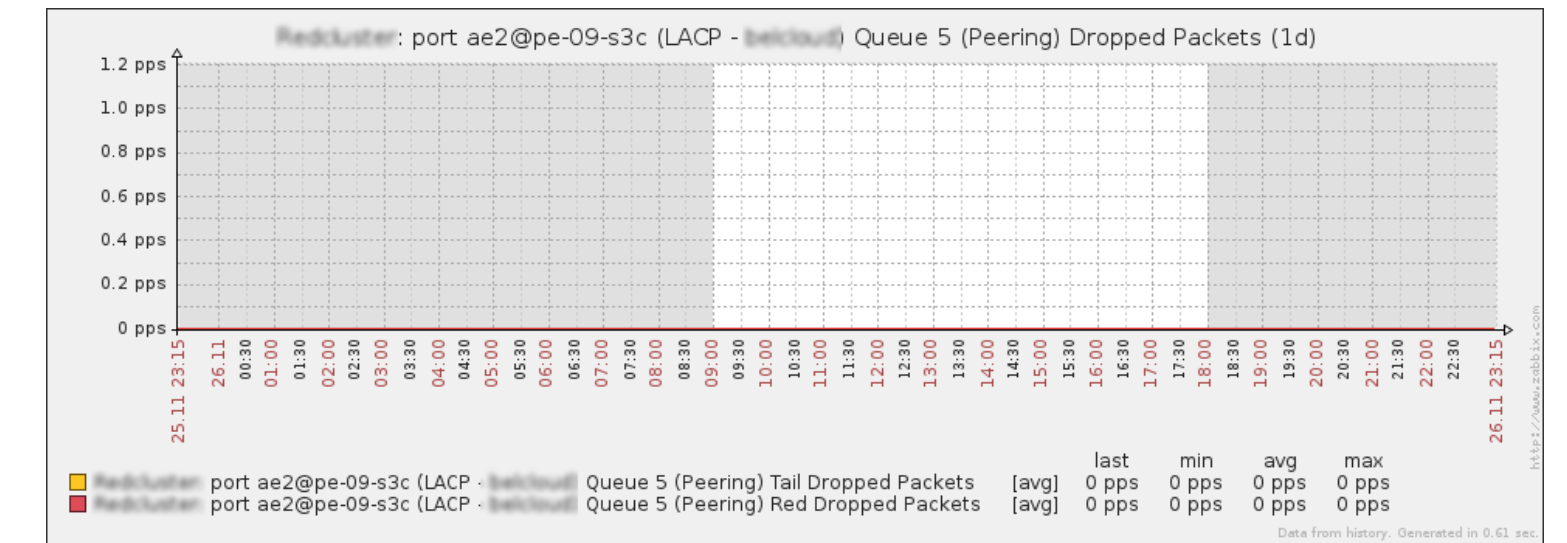
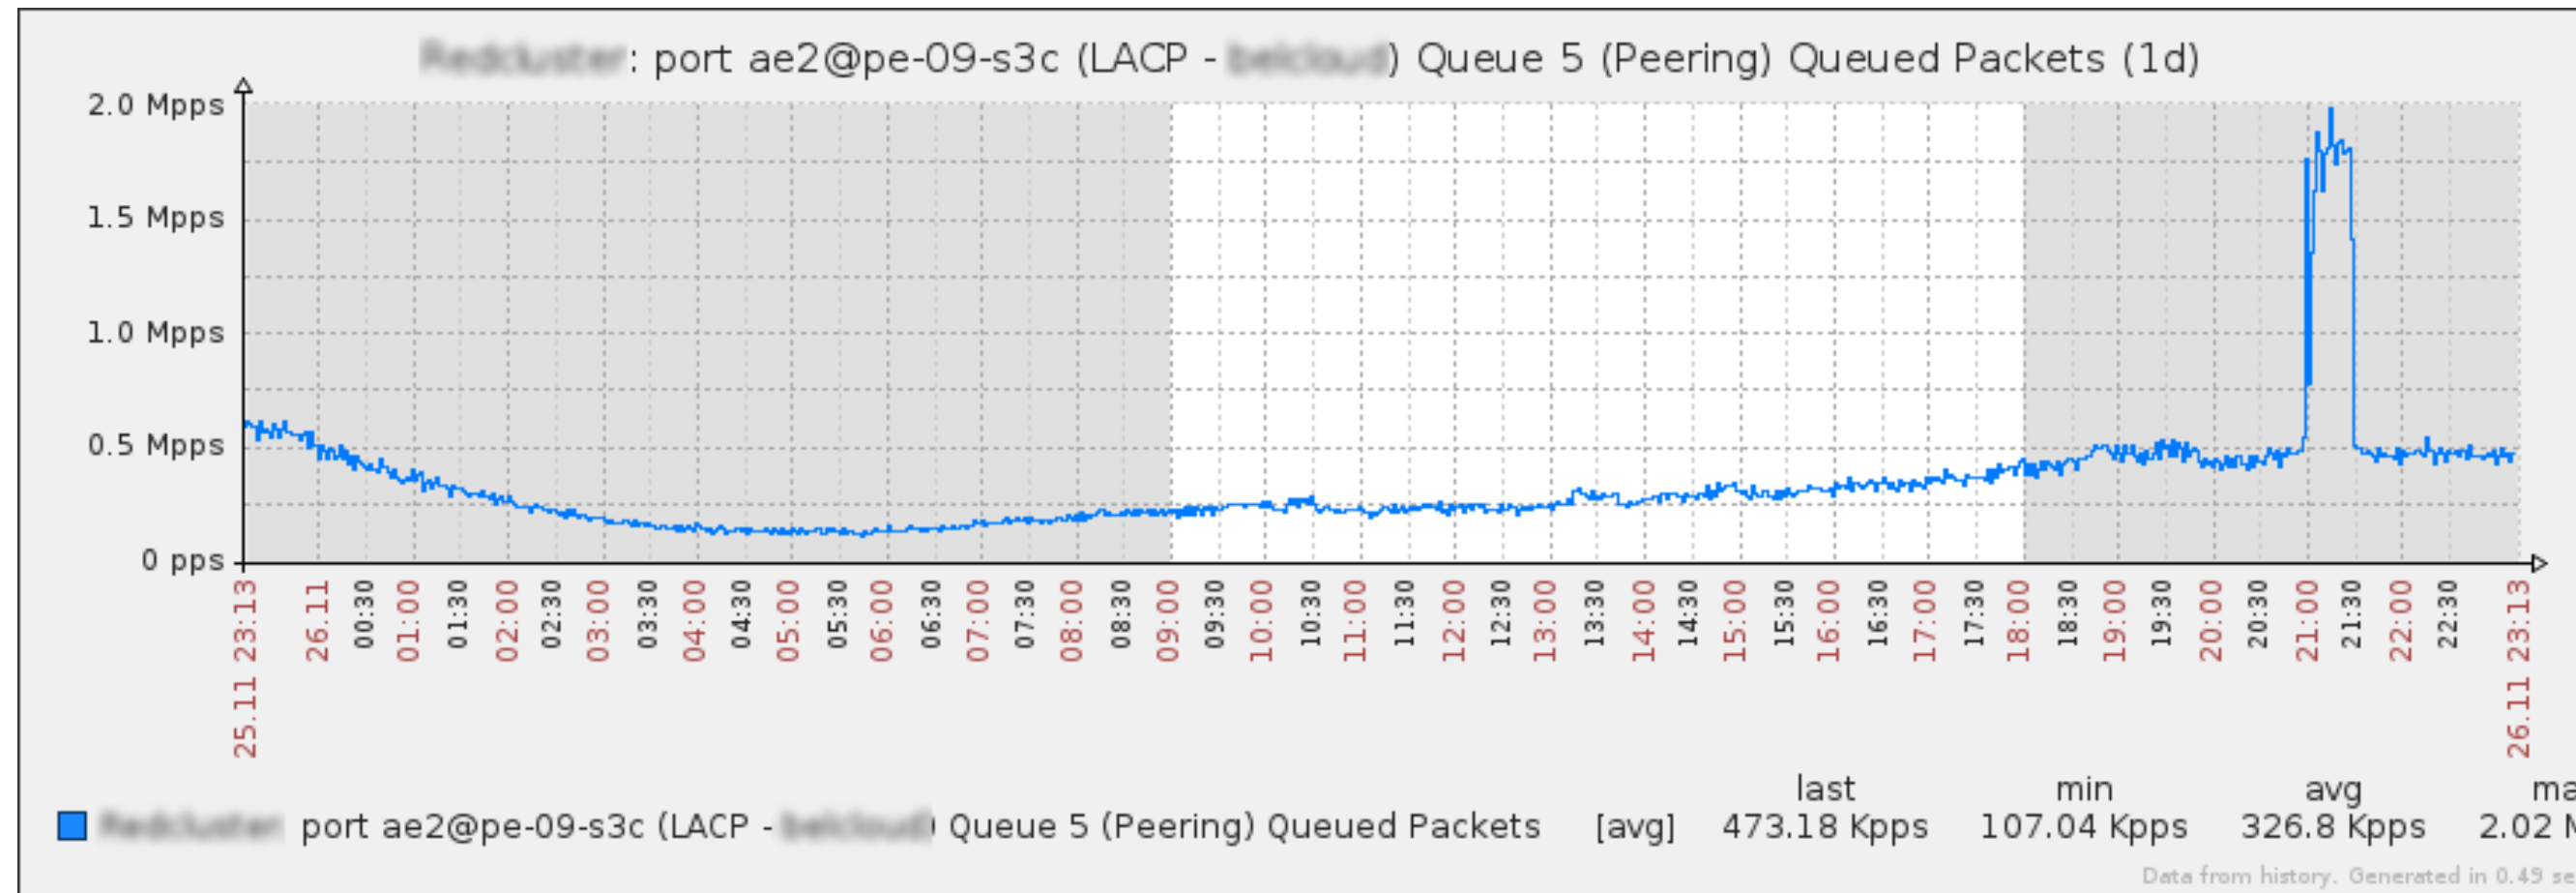
DDoS Queue Dropped in pps



DDoS Queue Queued in pps




Real Attack in Monitoring



Peering Queue Dropped in pps - 0

Peering Queue Queued in pps



Peering Queue traffic in Gbps








DDoS Solution





Pros & Cons

DDoS solution pros

-  Always on, works for shot time frame attacks
-  Whole network protected, including core links
-  Dot1p mark on output traffic







DDoS solution cons

-  Minimal possibility for classifying useful traffic as DDoS
-  Zero-day DDoS Attacks
-  Output Queue on Reseller ports
(Applied on port not VLAN)
-  Remote Peering Ports with lower speed than port speed



DDoS solution future development Ideas

-  Whitelisting
-  Calculate IP Address reputation
-  Per-VLAN queueing on Carrier/Reseller ports
-  Port speed shaping (Member configured in my.bix.bg)



Thank you!

Q&A