

A network diagram on the left side of the slide, consisting of several red circular nodes connected by thin grey lines. The nodes are arranged in a roughly triangular pattern, with some nodes having multiple connections to other nodes, creating a complex web of links.

How we built sFlow visualization tool (open source)

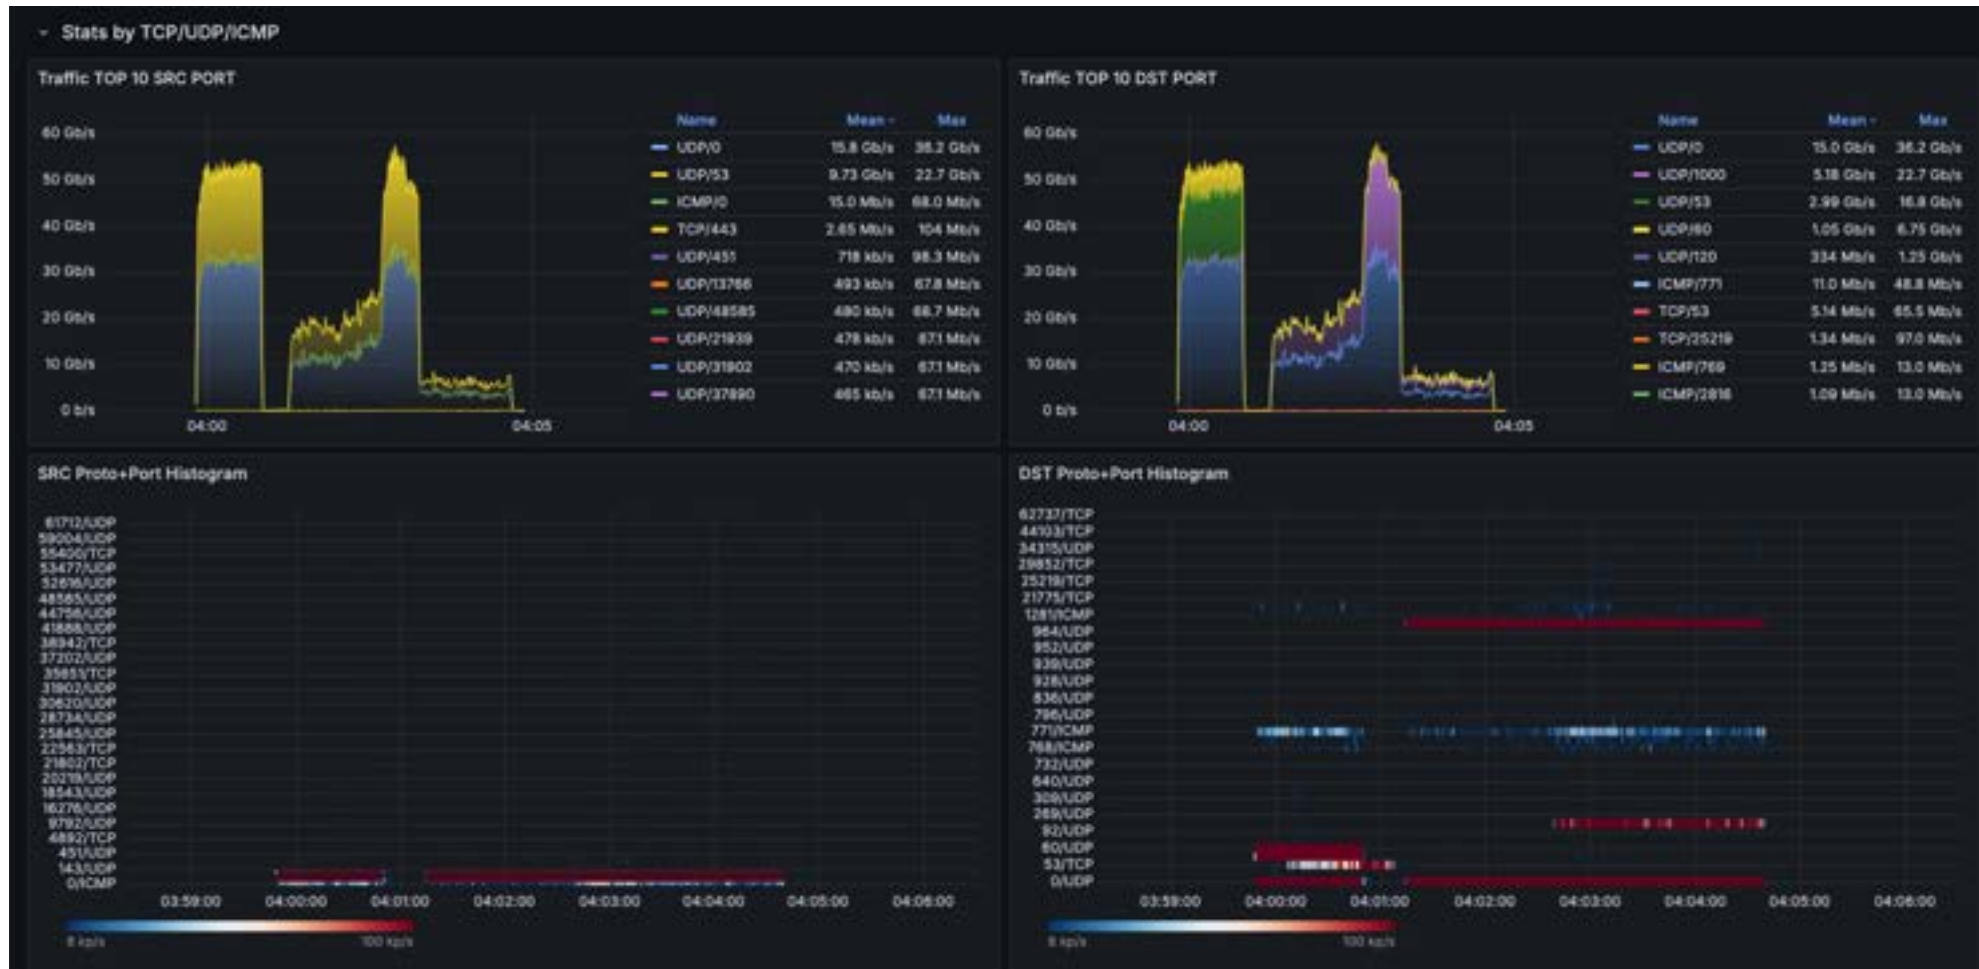
Blažej Krajňák
CSNOG | 22.1.2025

Energotel, a.s.

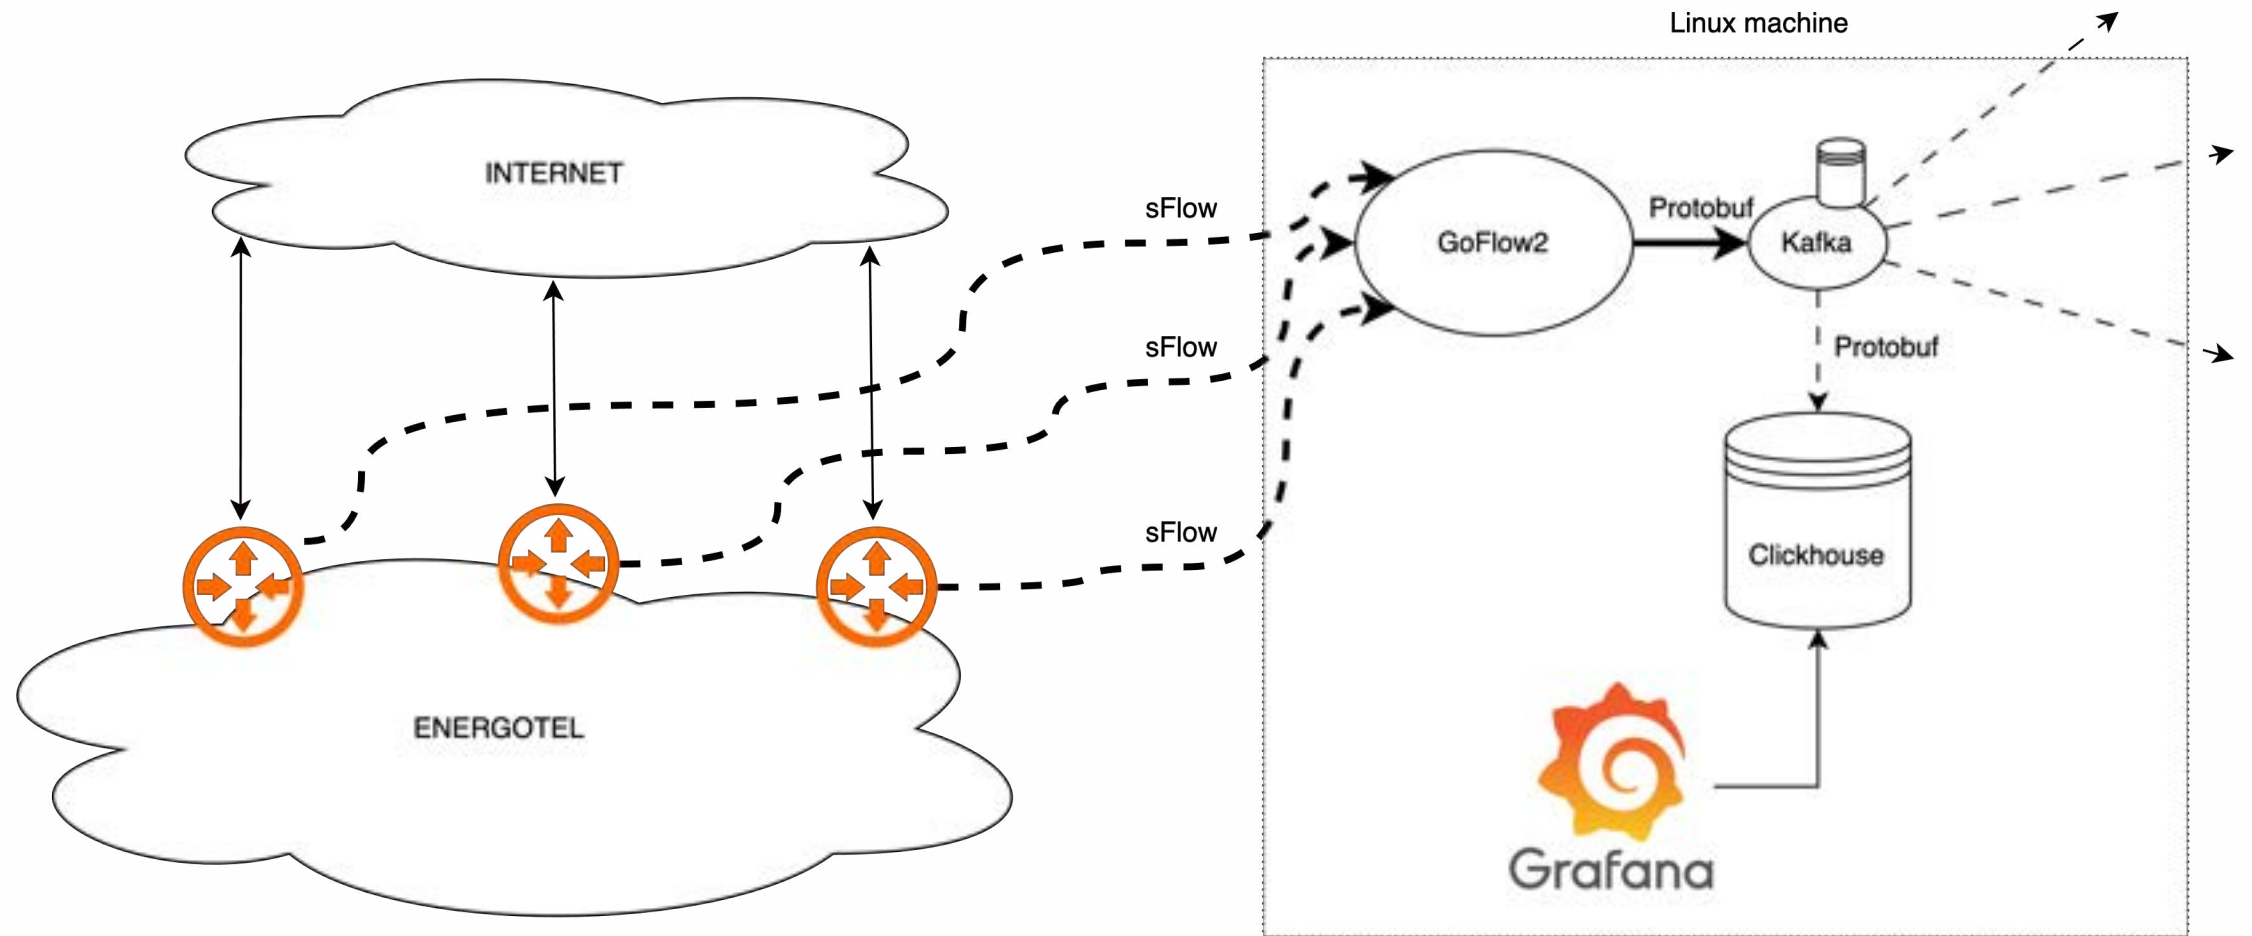
- provides B2B services
- national optical fiber network (400G backbone)
- internet services, L2 / L3 services
- lease of dark fibers / lambdas
- IT outsourcing
- 24/7 NOC and SOC
- cyber security audits and trainings
- part of critical national infrastructure



This comes soon ...



... but theory first – how the tool works



Clickhouse

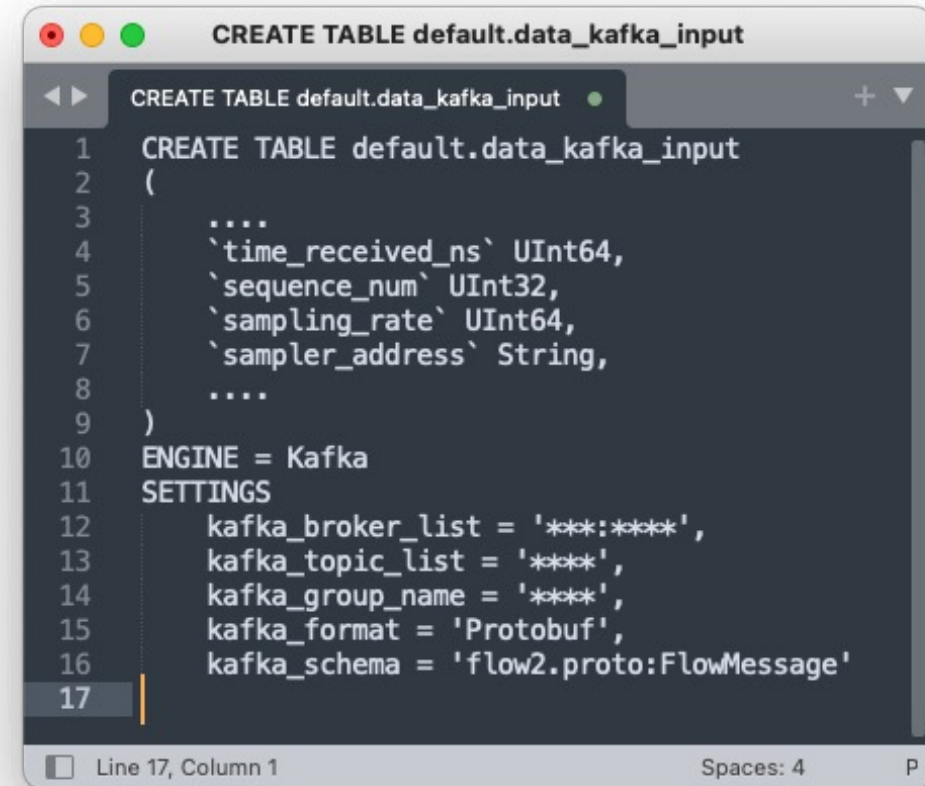


- ❑ column-oriented database management system
- ❑ data compression by default
- ❑ provides automatic data aggregation (materialized view)

Provides many nice features:

Clickhouse – Kafka connector

- reads and saves data from Kafka
- auto-parse protobuf (binary) data to structure



```
CREATE TABLE default.data_kafka_input
(
  ....
  `time_received_ns` UInt64,
  `sequence_num` UInt32,
  `sampling_rate` UInt64,
  `sampler_address` String,
  ....
)
ENGINE = Kafka
SETTINGS
  kafka_broker_list = '***:***',
  kafka_topic_list = '***',
  kafka_group_name = '***',
  kafka_format = 'Protobuf',
  kafka_schema = 'flow2.proto:FlowMessage'
```

Clickhouse – create table from online data

```
CREATE TABLE default.iana_protocols
(
  `Decimal` String,
  `Keyword` String,
  `Protocol` String,
  `IPv6ExtensionHeader` String,
  `Reference` String
)
ENGINE = URL('https://www.iana.org/assignments/protocol-numbers/protocol-numbers-1.csv', 'CSV')
SETTINGS input_format_csv_skip_first_lines = 1
```

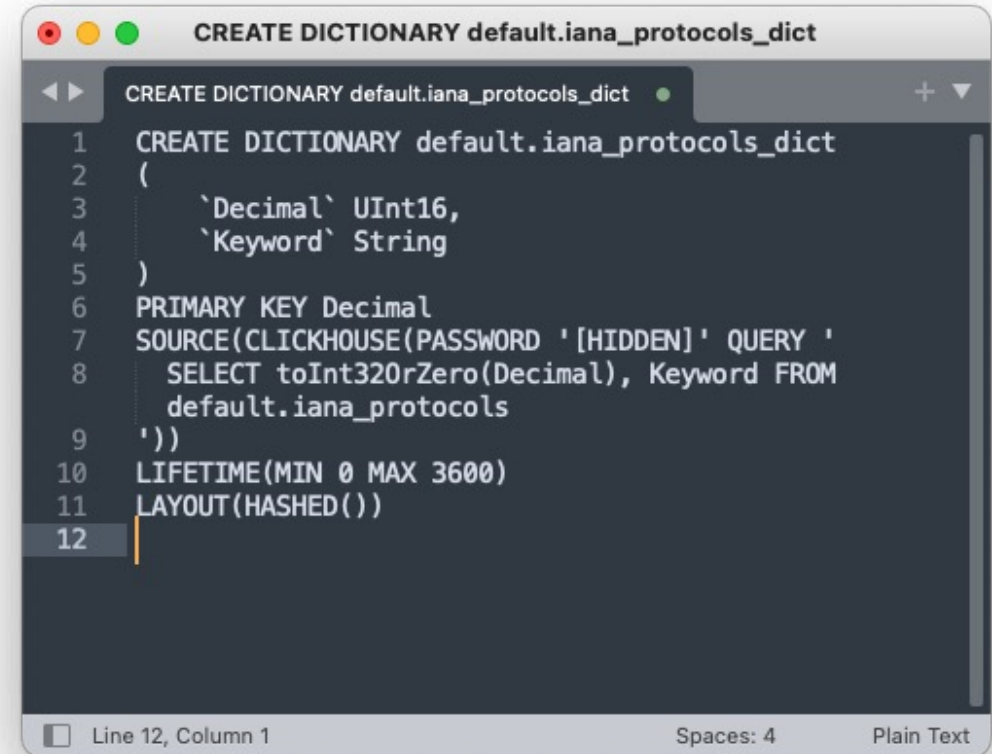
Line 11, Column 1 Spaces: 4 Plain Text

```
CREATE TABLE default.geoip
(
  `ip_range_start` IPv4,
  `ip_range_end` IPv4,
  `country_code` Nullable(String),
  `state1` Nullable(String),
  `state2` Nullable(String),
  `city` Nullable(String),
  `postcode` Nullable(String),
  `latitude` Float64,
  `longitude` Float64,
  `timezone` Nullable(String)
)
ENGINE = URL('https://raw.githubusercontent.com/sapics/ip-location-db/master/dbip-city/dbip-city-ipv4.csv.gz', 'CSV')
```

Line 15, Column 1 Spaces: 4 Plain Text

Clickhouse – dictionaries

- small tables kept in RAM



```
CREATE DICTIONARY default.iana_protocols_dict
(
  `Decimal` UInt16,
  `Keyword` String
)
PRIMARY KEY Decimal
SOURCE(CLICKHOUSE(PASSWORD '[HIDDEN]' QUERY '
SELECT toInt32orZero(Decimal), Keyword FROM
default.iana_protocols
'))
LIFETIME(MIN 0 MAX 3600)
LAYOUT(HASHED())
```

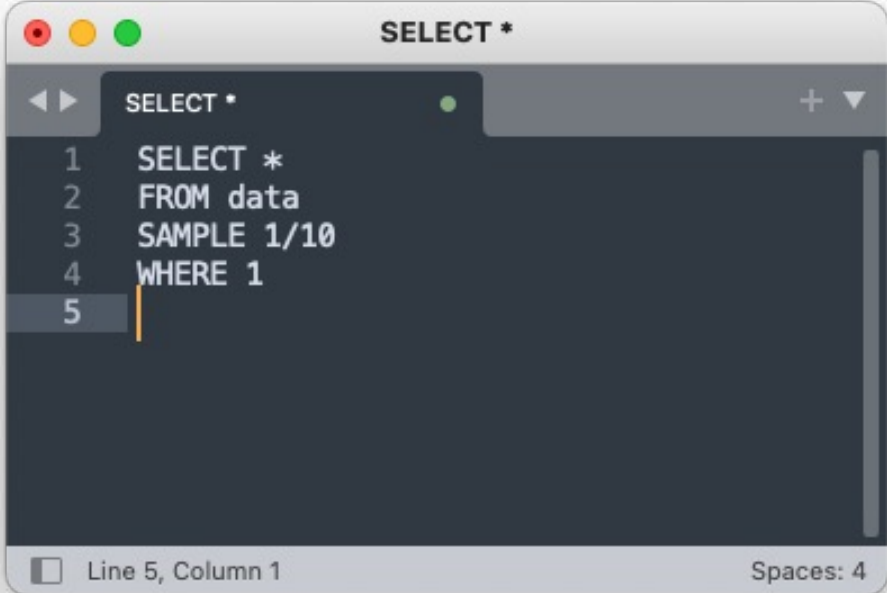
Line 12, Column 1 Spaces: 4 Plain Text

Clickhouse – ALIAS data type

```
CREATE TABLE default.data
(
    ....
    'ip_flags' UInt32,
    'tcp_flags' UInt32,
    ....
    'tcp_flag_fin' Bool ALIAS bitTest(tcp_flags, 0),
    'tcp_flag_syn' Bool ALIAS bitTest(tcp_flags, 1),
    'tcp_flag_rst' Bool ALIAS bitTest(tcp_flags, 2),
    'tcp_flag_psh' Bool ALIAS bitTest(tcp_flags, 3),
    'tcp_flag_ack' Bool ALIAS bitTest(tcp_flags, 4),
    'tcp_flag_urg' Bool ALIAS bitTest(tcp_flags, 5),
    'tcp_flag_ecn_echo' Bool ALIAS bitTest(tcp_flags, 6),
    'tcp_flag_cw_reduced' Bool ALIAS bitTest(tcp_flags, 7),
    'tcp_flag_acc_ecn' Bool ALIAS bitTest(tcp_flags, 8),
    ....
    'ip_flag_more_fragments' Bool ALIAS bitTest(ip_flags, 0),
    'ip_flag_dont_fragment' Bool ALIAS bitTest(ip_flags, 1),
    'ip_flag_reserved_bit' Bool ALIAS bitTest(ip_flags, 2),
    ....
    'src_as_name' String ALIAS concat('AS', src_as, ' - ', dictGetOrDefault(default.asn_dict, 'asn_name', src_as, 'Unknown')),
    'dst_as_name' String ALIAS concat('AS', dst_as, ' - ', dictGetOrDefault(default.asn_dict, 'asn_name', dst_as, 'Unknown')),
    'in_if_name' String ALIAS dictGetOrDefault(default.engt_ifaces_dict, 'ifDescription', (sampler_ip, in_if), concat(sampler_ip, '-', in_if)),
    'out_if_name' String ALIAS dictGetOrDefault(default.engt_ifaces_dict, 'ifDescription', (sampler_ip, out_if), concat(sampler_ip, '-', out_if)),
    'etype_name' String ALIAS dictGetOrDefault(default.ethertypes_dict, 'Keyword', etype, etype),
    'proto_name' String ALIAS dictGetOrDefault(default.iana_protocols_dict, 'Keyword', proto, proto),
    'src_mac_name' String ALIAS MACNumToString(src_mac),
    'dst_mac_name' String ALIAS MACNumToString(dst_mac),
    'ip_dscp' UInt32 ALIAS bitShiftRight(ip_tos, 2),
    'dir_in' Bool ALIAS dictHas(default.engt_ifaces_dict, (sampler_ip, in_if)),
    'dir_out' Bool ALIAS dictHas(default.engt_ifaces_dict, (sampler_ip, out_if)),
    ....
)
```

Clickhouse – SAMPLE function

- performs query on fraction of data (1/n)
- helps to speed up query (where suitable)
- it's a deterministic mechanism
- need to specify sampling key on table creation

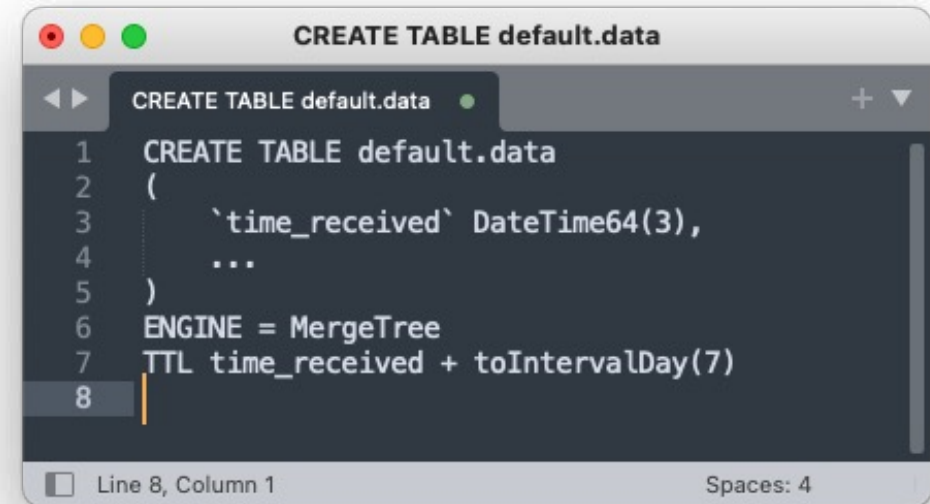


```
SELECT *
FROM data
SAMPLE 1/10
WHERE 1
```

The screenshot shows a SQL editor window titled "SELECT *". The query text is: "SELECT * FROM data SAMPLE 1/10 WHERE 1". The cursor is positioned at the end of line 5, column 1. The status bar at the bottom indicates "Line 5, Column 1" and "Spaces: 4".

Clickhouse – auto data removal

- ability to specify data lifetime

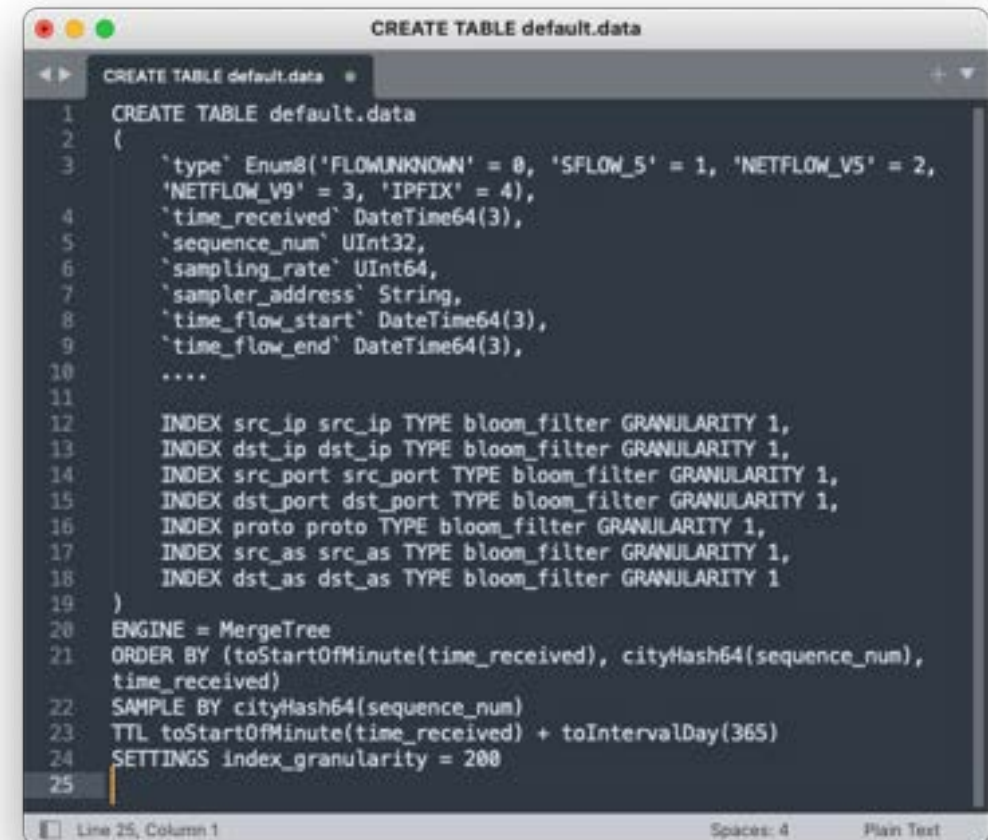


```
CREATE TABLE default.data
(
  `time_received` DateTime64(3),
  ...
)
ENGINE = MergeTree
TTL time_received + toIntervalDay(7)
```

The screenshot shows a code editor window titled "CREATE TABLE default.data". The code defines a table with a MergeTree engine and a TTL clause that sets the data to expire 7 days after the 'time_received' column value. The editor shows line numbers 1 through 8, with the cursor at the end of line 8.

Clickhouse – the main data table tuning

- sampling key needs to be random INT value → use hash
- sampling key needs to be present in primary key (= ORDER BY if not specified)
- filtering by time range is more important than sampling key → use „time_received“ before sampling key in primary key
- put only low-cardinality columns before sampling expression in primary key -> use toStartOfMinute(time_received)
- any INDEX in Clickhouse is skipping index -> use small enough granules size



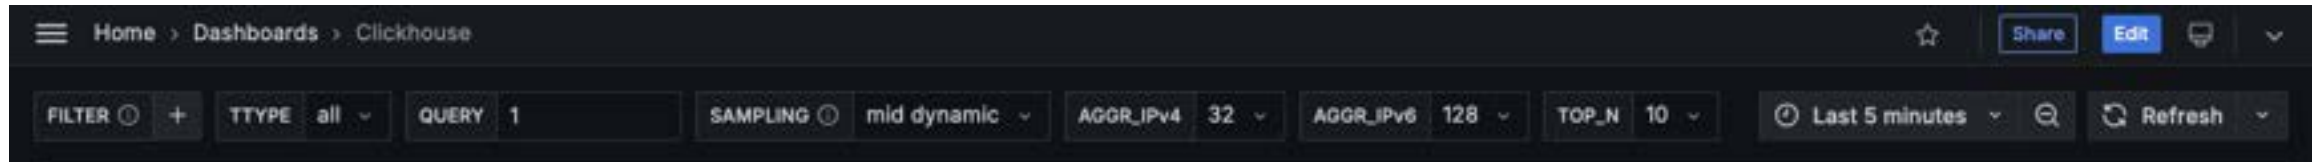
```
CREATE TABLE default.data
(
  `type` Enum8('FLOWUNKNOWN' = 0, 'SFLOW_5' = 1, 'NETFLOW_V5' = 2,
  'NETFLOW_V9' = 3, 'IPFIX' = 4),
  `time_received` DateTime64(3),
  `sequence_num` UInt32,
  `sampling_rate` UInt64,
  `sampler_address` String,
  `time_flow_start` DateTime64(3),
  `time_flow_end` DateTime64(3),
  ....
  INDEX src_ip src_ip TYPE bloom_filter GRANULARITY 1,
  INDEX dst_ip dst_ip TYPE bloom_filter GRANULARITY 1,
  INDEX src_port src_port TYPE bloom_filter GRANULARITY 1,
  INDEX dst_port dst_port TYPE bloom_filter GRANULARITY 1,
  INDEX proto proto TYPE bloom_filter GRANULARITY 1,
  INDEX src_as src_as TYPE bloom_filter GRANULARITY 1,
  INDEX dst_as dst_as TYPE bloom_filter GRANULARITY 1
)
ENGINE = MergeTree
ORDER BY (toStartOfMinute(time_received), cityHash64(sequence_num),
time_received)
SAMPLE BY cityHash64(sequence_num)
TTL toStartOfMinute(time_received) + toIntervalDay(365)
SETTINGS index_granularity = 200
```

Available data metrics (from GoFlow2 + computed)

```
`type` Enum8('FLOWUNKNOWN' = 0,
'SFLOW_5' = 1, 'NETFLOW_V5' = 2,
'NETFLOW_V9' = 3, 'IPFIX' = 4),
`time_received` DateTime64(3),
`sequence_num` UInt32,
`sampling_rate` UInt64,
`sampler_address` String,
`time_flow_start` DateTime64(3),
`time_flow_end` DateTime64(3),
`bytes` UInt64,
`packets` UInt64,
`src_addr` String,
`dst_addr` String,
`etype` UInt32,
`proto` UInt32,
`src_port` UInt32,
`dst_port` UInt32,
`in_if` UInt32,
`out_if` UInt32,
`src_mac` UInt64,
`dst_mac` UInt64,
`src_vlan` UInt32,
`dst_vlan` UInt32,
`vlan_id` UInt32,
`ip_tos` UInt32,
`forwarding_status` UInt32,
`ip_ttl` UInt32,
`ip_flags` UInt32,
`tcp_flags` UInt32,
`icmp_type` UInt32,
`icmp_code` UInt32,
`ipv6_flow_label` UInt32,
`fragment_id` UInt32,
`fragment_offset` UInt32,
`src_as` UInt32,
`dst_as` UInt32,
`next_hop` String,
`next_hop_as` UInt32,
`src_net` UInt32,
`dst_net` UInt32,
`bgp_next_hop` String,
`bgp_communities` Array(UInt32),
`as_path` Array(UInt32),
`mpls_ttl` Array(UInt32),
`mpls_label` Array(UInt32),
`mpls_ip` Array(String),
`observation_domain_id` UInt32,
`observation_point_id` UInt32,
`sampler_ip` IPv6,
`src_ip` IPv6,
`dst_ip` IPv6,
`bgp_next_hop_ip` IPv6,
`tcp_flag_fin` Bool ALIAS bitTest(tcp_flags, 0),
`tcp_flag_syn` Bool ALIAS bitTest(tcp_flags, 1),
`tcp_flag_rst` Bool ALIAS bitTest(tcp_flags, 2),
`tcp_flag_psh` Bool ALIAS bitTest(tcp_flags, 3),
`tcp_flag_ack` Bool ALIAS bitTest(tcp_flags, 4),
`tcp_flag_urg` Bool ALIAS bitTest(tcp_flags, 5),
`tcp_flag_ecn_echo` Bool ALIAS bitTest(tcp_flags, 6),
`tcp_flag_cw_reduced` Bool ALIAS bitTest(tcp_flags, 7),
`tcp_flag_acc_ecn` Bool ALIAS bitTest(tcp_flags, 8),
`ip_flag_more_fragments` Bool ALIAS bitTest(ip_flags, 0),
`ip_flag_dont_fragment` Bool ALIAS bitTest(ip_flags, 1),
`ip_flag_reserved_bit` Bool ALIAS bitTest(ip_flags, 2),
`src_as_name` String ALIAS concat('AS', src_as, ' - ', dictGetOrDefault(default.asn_dict, 'asn_name', src_as, 'Unknown')),
`dst_as_name` String ALIAS concat('AS', dst_as, ' - ', dictGetOrDefault(default.asn_dict, 'asn_name', dst_as, 'Unknown')),
`in_if_name` String ALIAS dictGetOrDefault(default.engt_ifaces_dict, 'ifDescription', (sampler_ip, in_if),
concat(sampler_ip, '-', in_if)),
`out_if_name` String ALIAS dictGetOrDefault(default.engt_ifaces_dict, 'ifDescription', (sampler_ip, out_if),
concat(sampler_ip, '-', out_if)),
`etype_name` String ALIAS dictGetOrDefault(default.ethertypes_dict, 'Keyword', etype, etype),
`proto_name` String ALIAS dictGetOrDefault(default.iana_protocols_dict, 'Keyword', proto, proto),
`src_mac_name` String ALIAS MACNumToString(src_mac),
`dst_mac_name` String ALIAS MACNumToString(dst_mac),
`ip_dscp` UInt32 ALIAS bitShiftRight(ip_tos, 2),
`dir_in` Bool ALIAS dictHas(default.engt_ifaces_dict, (sampler_ip, in_if)),
`dir_out` Bool ALIAS dictHas(default.engt_ifaces_dict, (sampler_ip, out_if))
```

Grafana dashboard

Dashboard controls

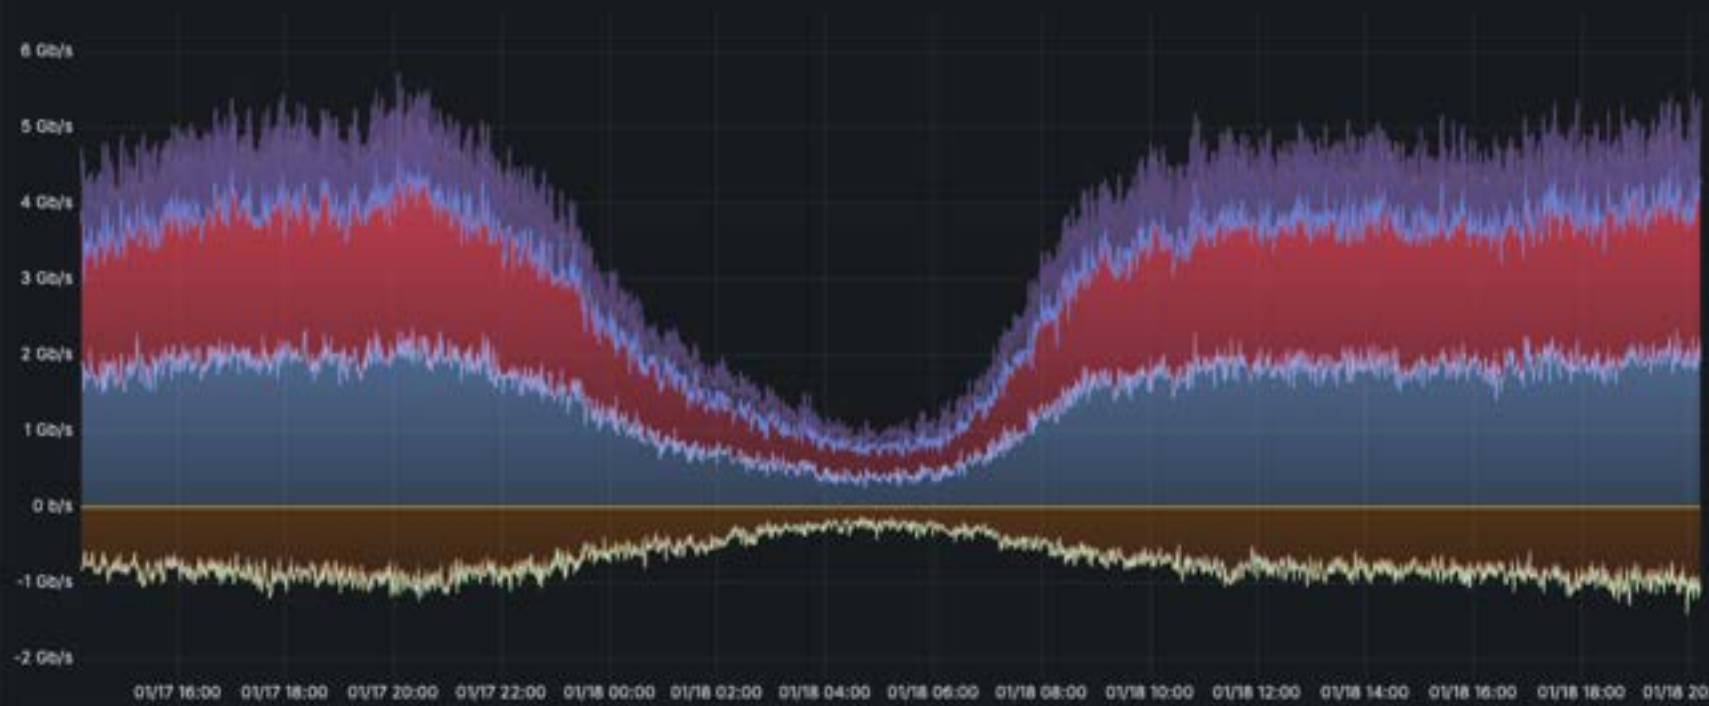


- ❑ FILTER – „user frendy“ filter (with autocomplete)
- ❑ TTYTYPE – predefined filters (SYN only, fragments, UDP 53/161, ...)
- ❑ QUERY – custom SQL filter statements (not colleague friendly but my favorite)
- ❑ SAMPLING – static or dynamic sampling rate when generating graphs (dynamic adjusts by actual time window - shorter the window, lower the sampling rate till 1/1)
- ❑ AGGR_IPv4 / AGGR_IPv6 – cumulate data by specified IP range
- ❑ TOP_N – how many TOP N items to render in graphs in whole dashboard



Stats by IN / OUT Interface

Traffic bps by Interface



Name	Mean	Max	Last
NIX.SK IN	1.54 Gb/s	2.38 Gb/s	2.34 Gb/s
NIX.CZ IN	1.53 Gb/s	2.37 Gb/s	1.92 Gb/s
PEERING.CZ IN	667 Mb/s	1.18 Gb/s	951 Mb/s
ORANGE IN	8.87 Mb/s	58.6 Mb/s	31.4 Mb/s
SIX.SK IN	5.28 Mb/s	62.8 Mb/s	262 kb/s
SUPERNET_02 IN	4.99 Mb/s	49.6 Mb/s	655 kb/s
SUPERNET_01 IN	4.75 Mb/s	41.7 Mb/s	0 b/s
SWAN.SK IN	288 kb/s	23.4 Mb/s	0 b/s
SIX2.SK IN	45.5 kb/s	9.52 Mb/s	0 b/s
1-IX IN	1.03 kb/s	465 kb/s	0 b/s
1-IX OUT	-237 b/s	0 b/s	0 b/s
SUPERNET_02 OUT	-2.54 Mb/s	0 b/s	0 b/s
SWAN.SK OUT	-3.17 Mb/s	0 b/s	-629 kb/s
SIX2.SK OUT	-3.49 Mb/s	0 b/s	0 b/s
ORANGE OUT	-13.2 Mb/s	0 b/s	-10.4 Mb/s
PEERING.CZ OUT	-14.6 Mb/s	0 b/s	-25.5 Mb/s
NIX.CZ OUT	-686 Mb/s	-108 Mb/s	-1.08 Gb/s

IP Analysis

EtherTypes TOP 10



IP Flags



IP Analysis

EtherTypes TOP 10



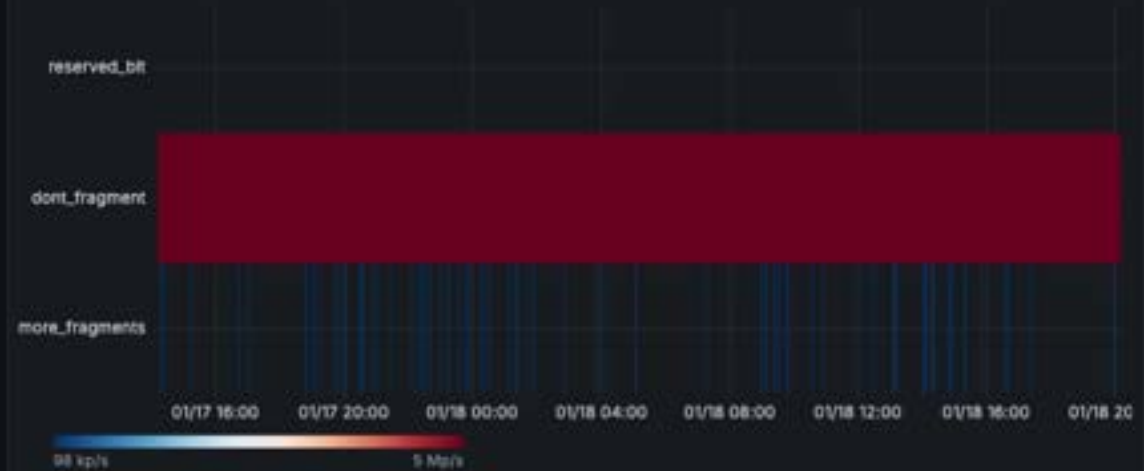
IP Flags



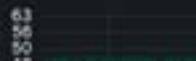
Fragment Offset Histogram



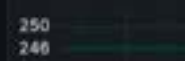
IP Flags Histogram

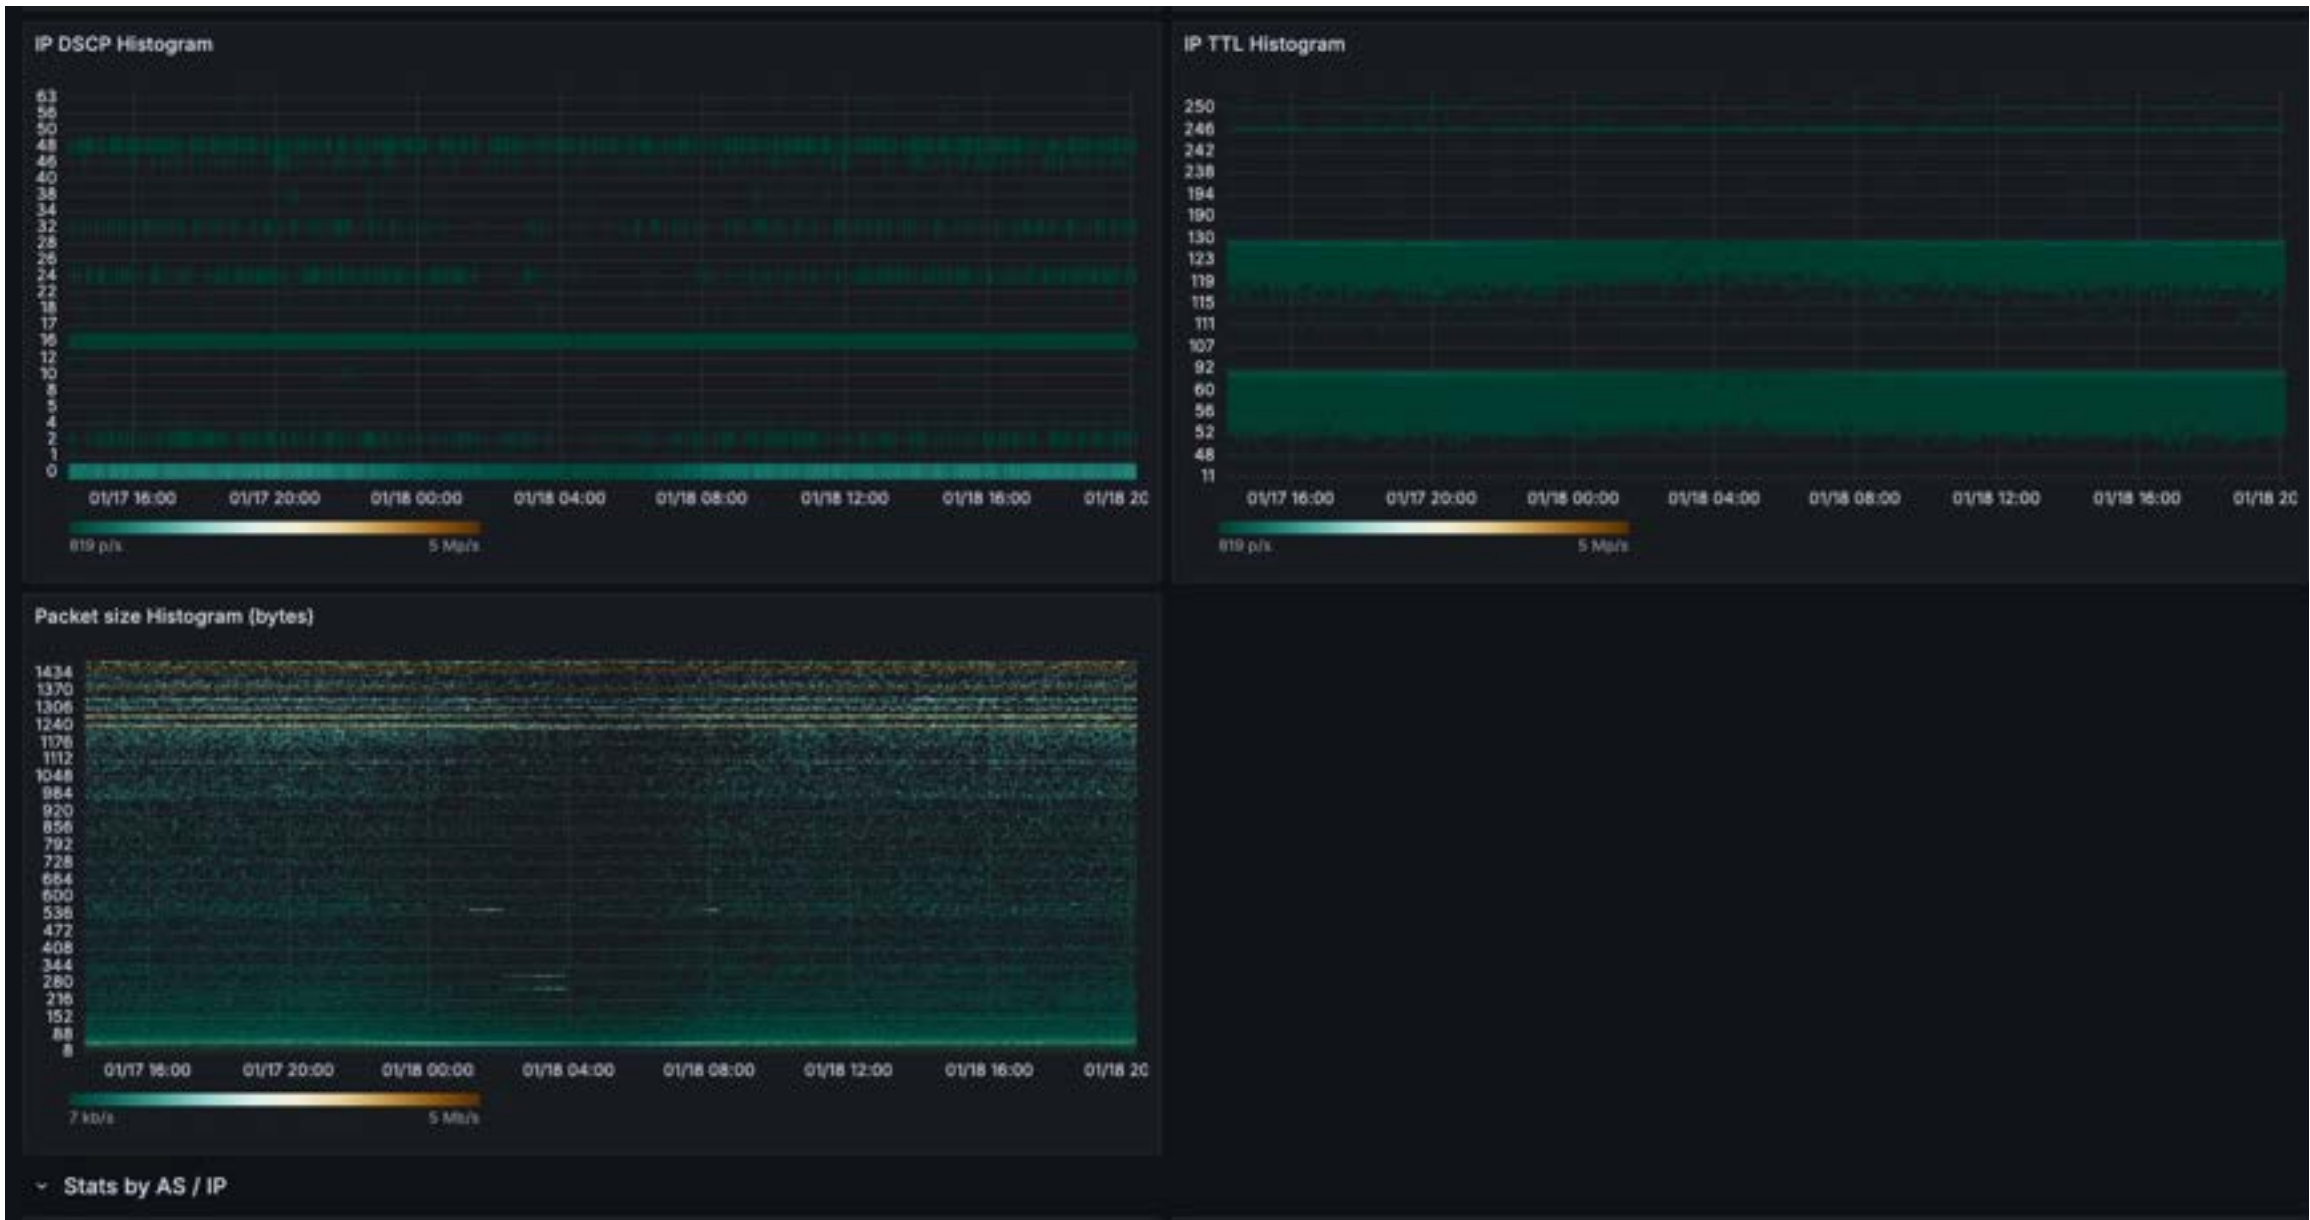


IP DSCP Histogram



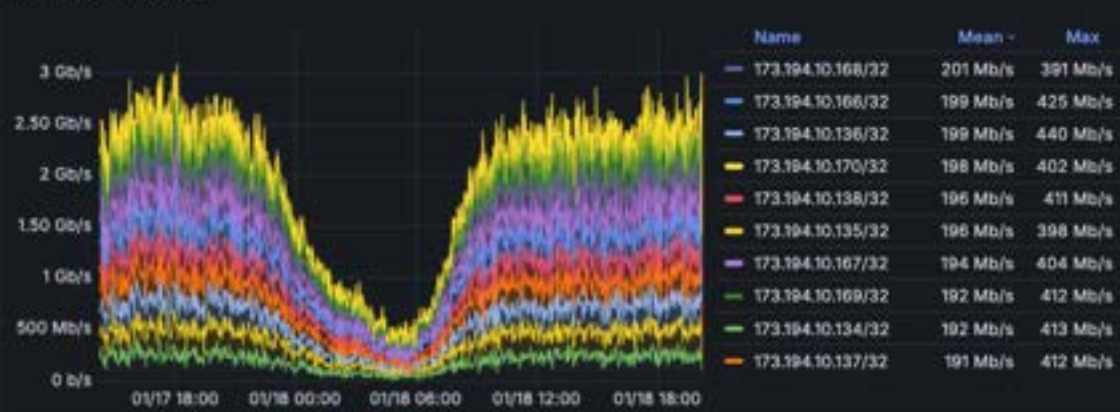
IP TTL Histogram





Stats by AS / IP

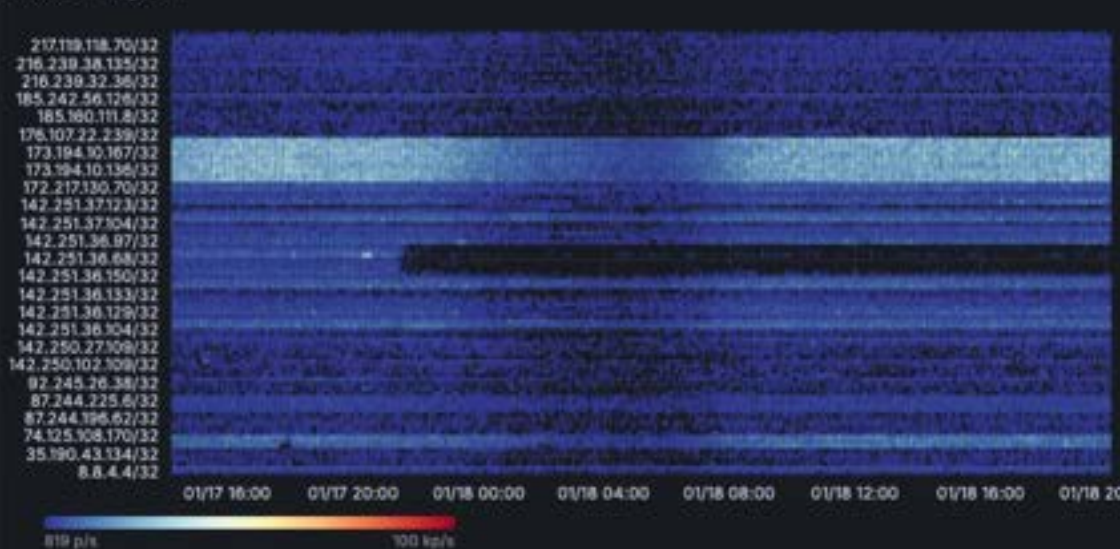
Traffic TOP 10 SRC IP



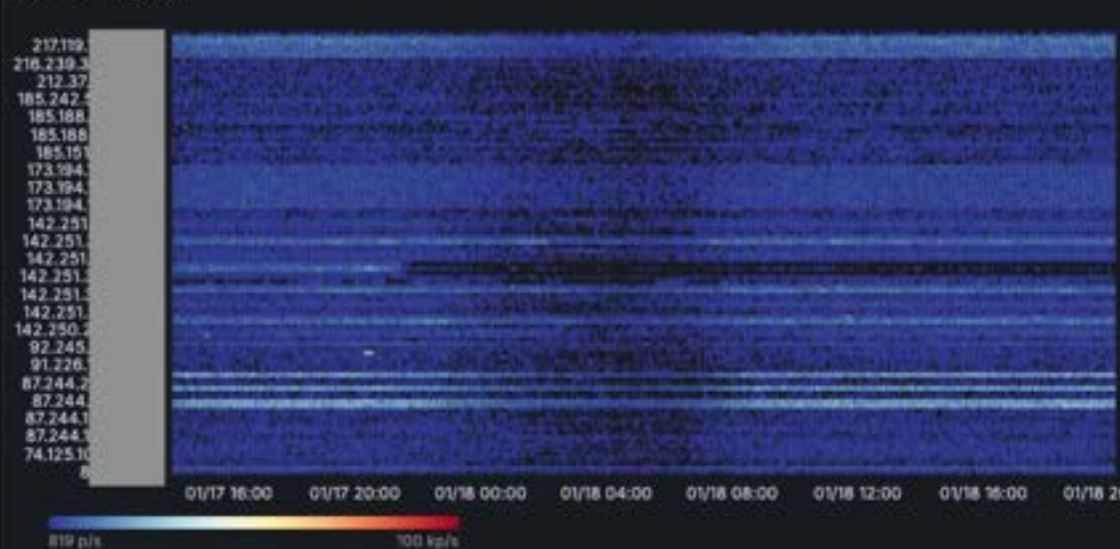
Traffic TOP 10 DST IP



SRC IP Histogram

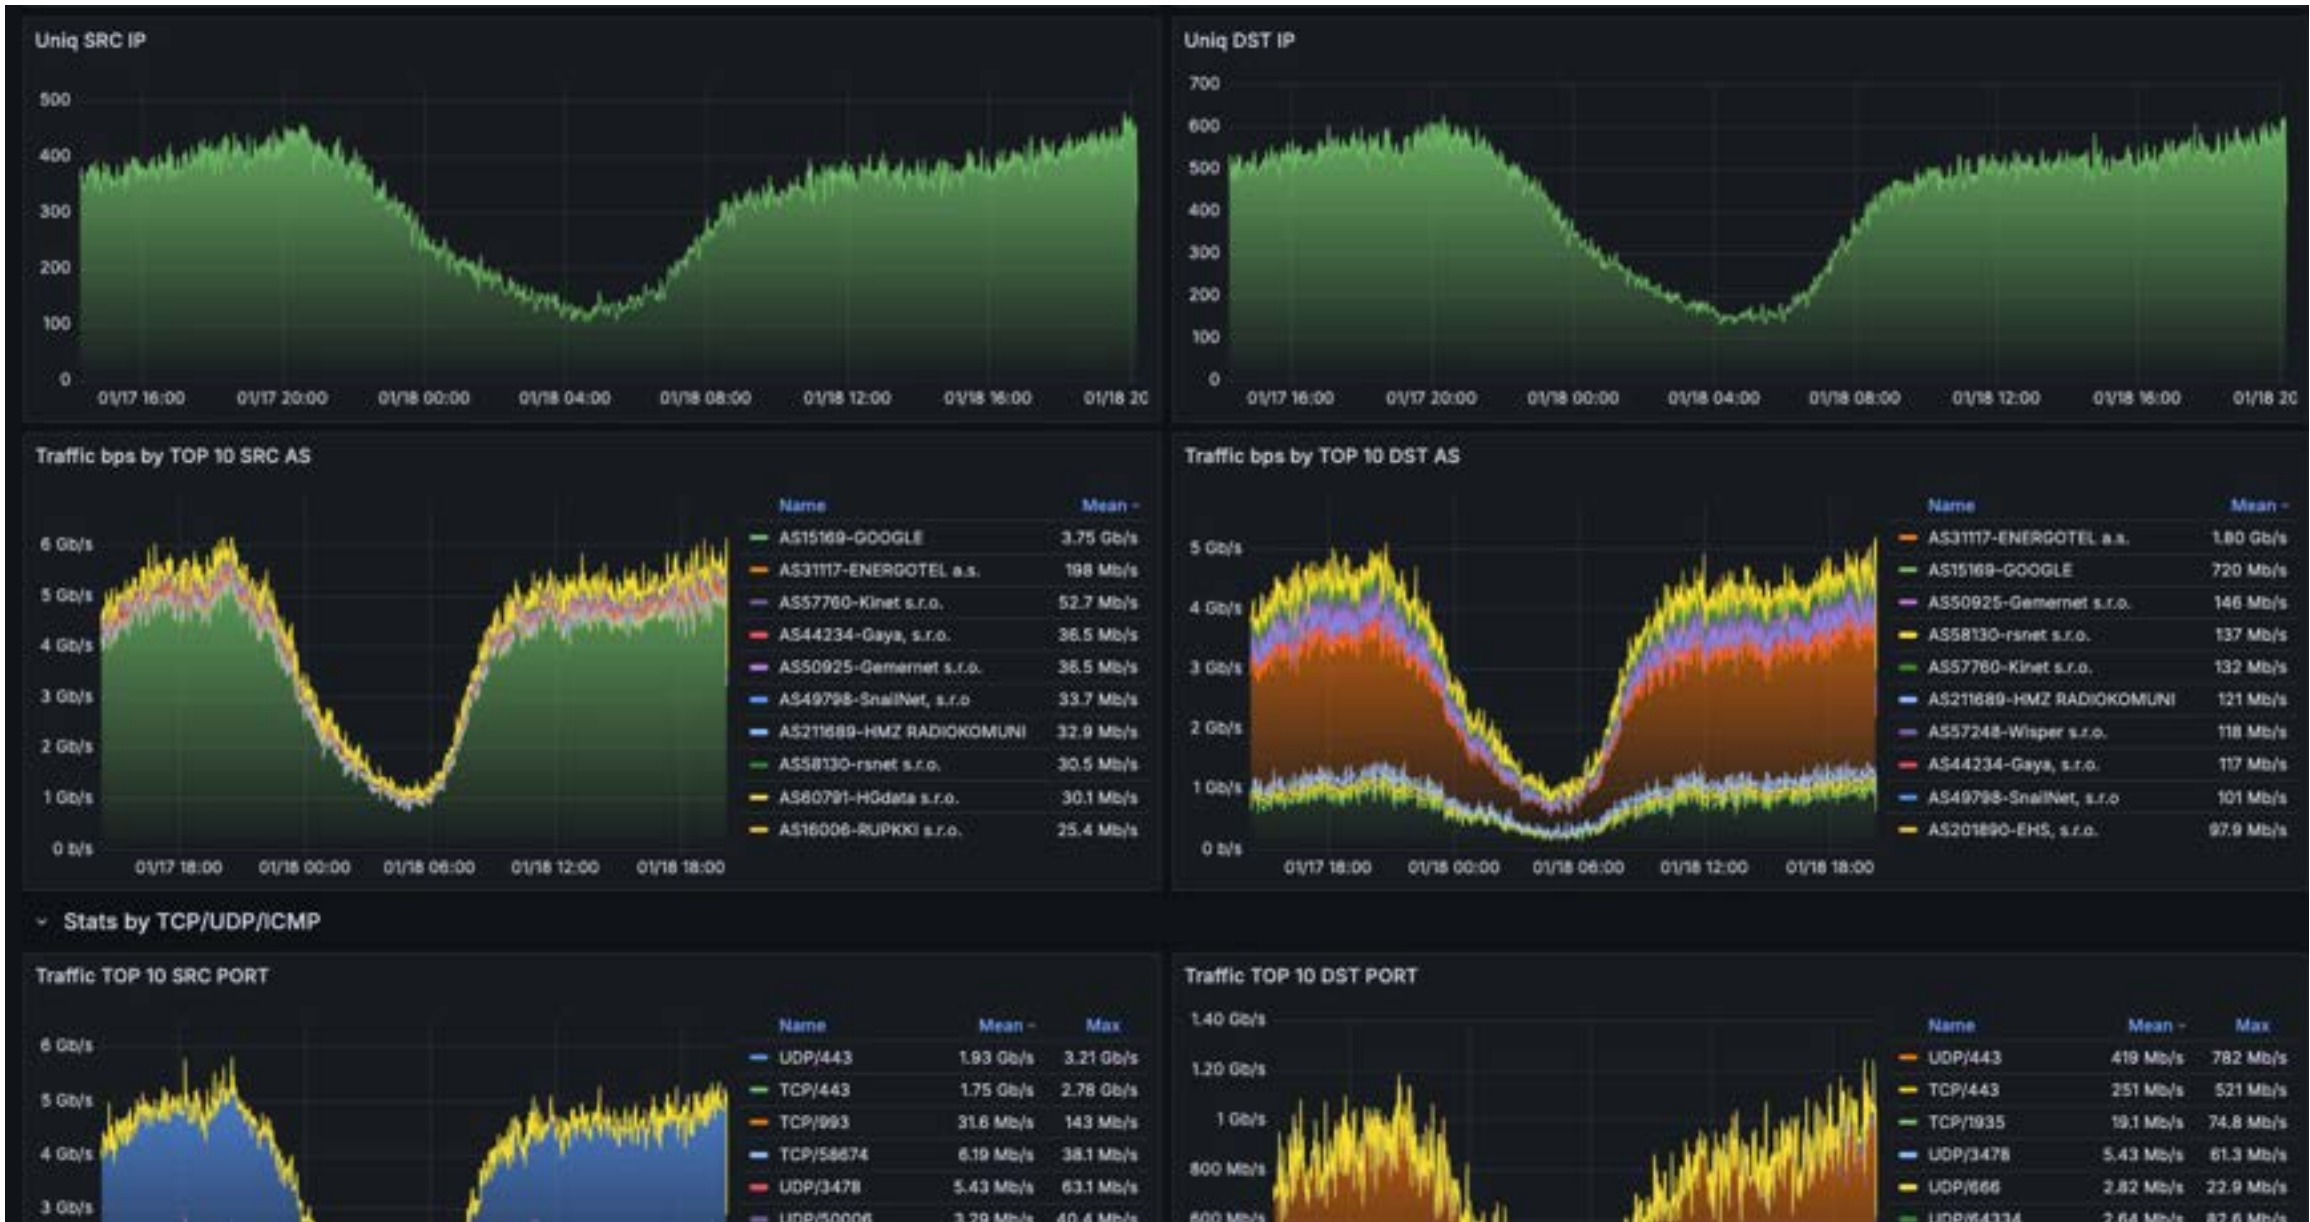


DST IP Histogram



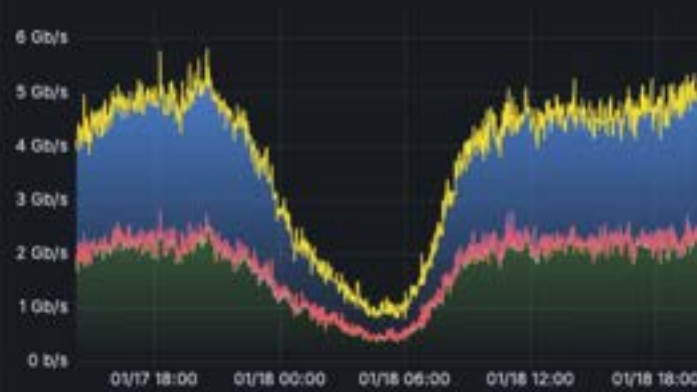
Uniq SRC IP

Uniq DST IP



Stats by TCP/UDP/ICMP

Traffic TOP 10 SRC PORT



Name	Mean	Max
UDP/443	1.93 Gb/s	3.21 Gb/s
TCP/443	1.75 Gb/s	2.78 Gb/s
TCP/993	31.6 Mb/s	143 Mb/s
TCP/58674	6.19 Mb/s	38.1 Mb/s
UDP/3478	5.43 Mb/s	63.1 Mb/s
UDP/50006	3.29 Mb/s	40.4 Mb/s
UDP/53	3.21 Mb/s	16.6 Mb/s
UDP/50008	2.93 Mb/s	34.9 Mb/s
UDP/50005	2.84 Mb/s	25.7 Mb/s
TCP/5228	2.70 Mb/s	27.4 Mb/s

Traffic TOP 10 DST PORT

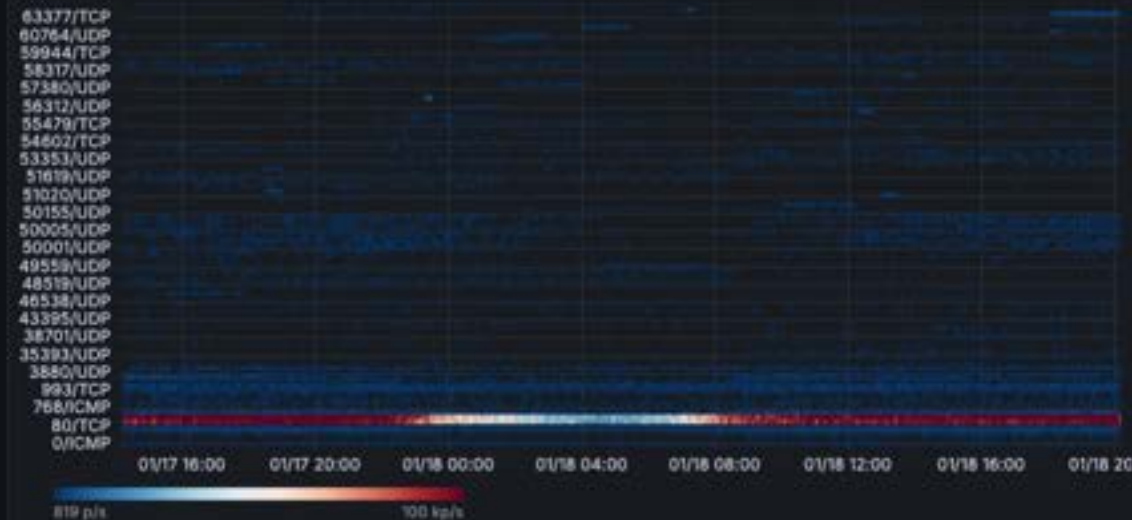


Name	Mean	Max
UDP/443	419 Mb/s	782 Mb/s
TCP/443	251 Mb/s	521 Mb/s
TCP/1935	19.1 Mb/s	74.8 Mb/s
UDP/3478	5.43 Mb/s	61.3 Mb/s
UDP/666	2.82 Mb/s	22.9 Mb/s
UDP/84334	2.64 Mb/s	82.6 Mb/s
UDP/50005	2.10 Mb/s	31.5 Mb/s
UDP/50008	1.99 Mb/s	32.4 Mb/s
UDP/50003	1.89 Mb/s	31.1 Mb/s
UDP/50004	1.40 Mb/s	23.8 Mb/s

SRC Proto+Port Histogram

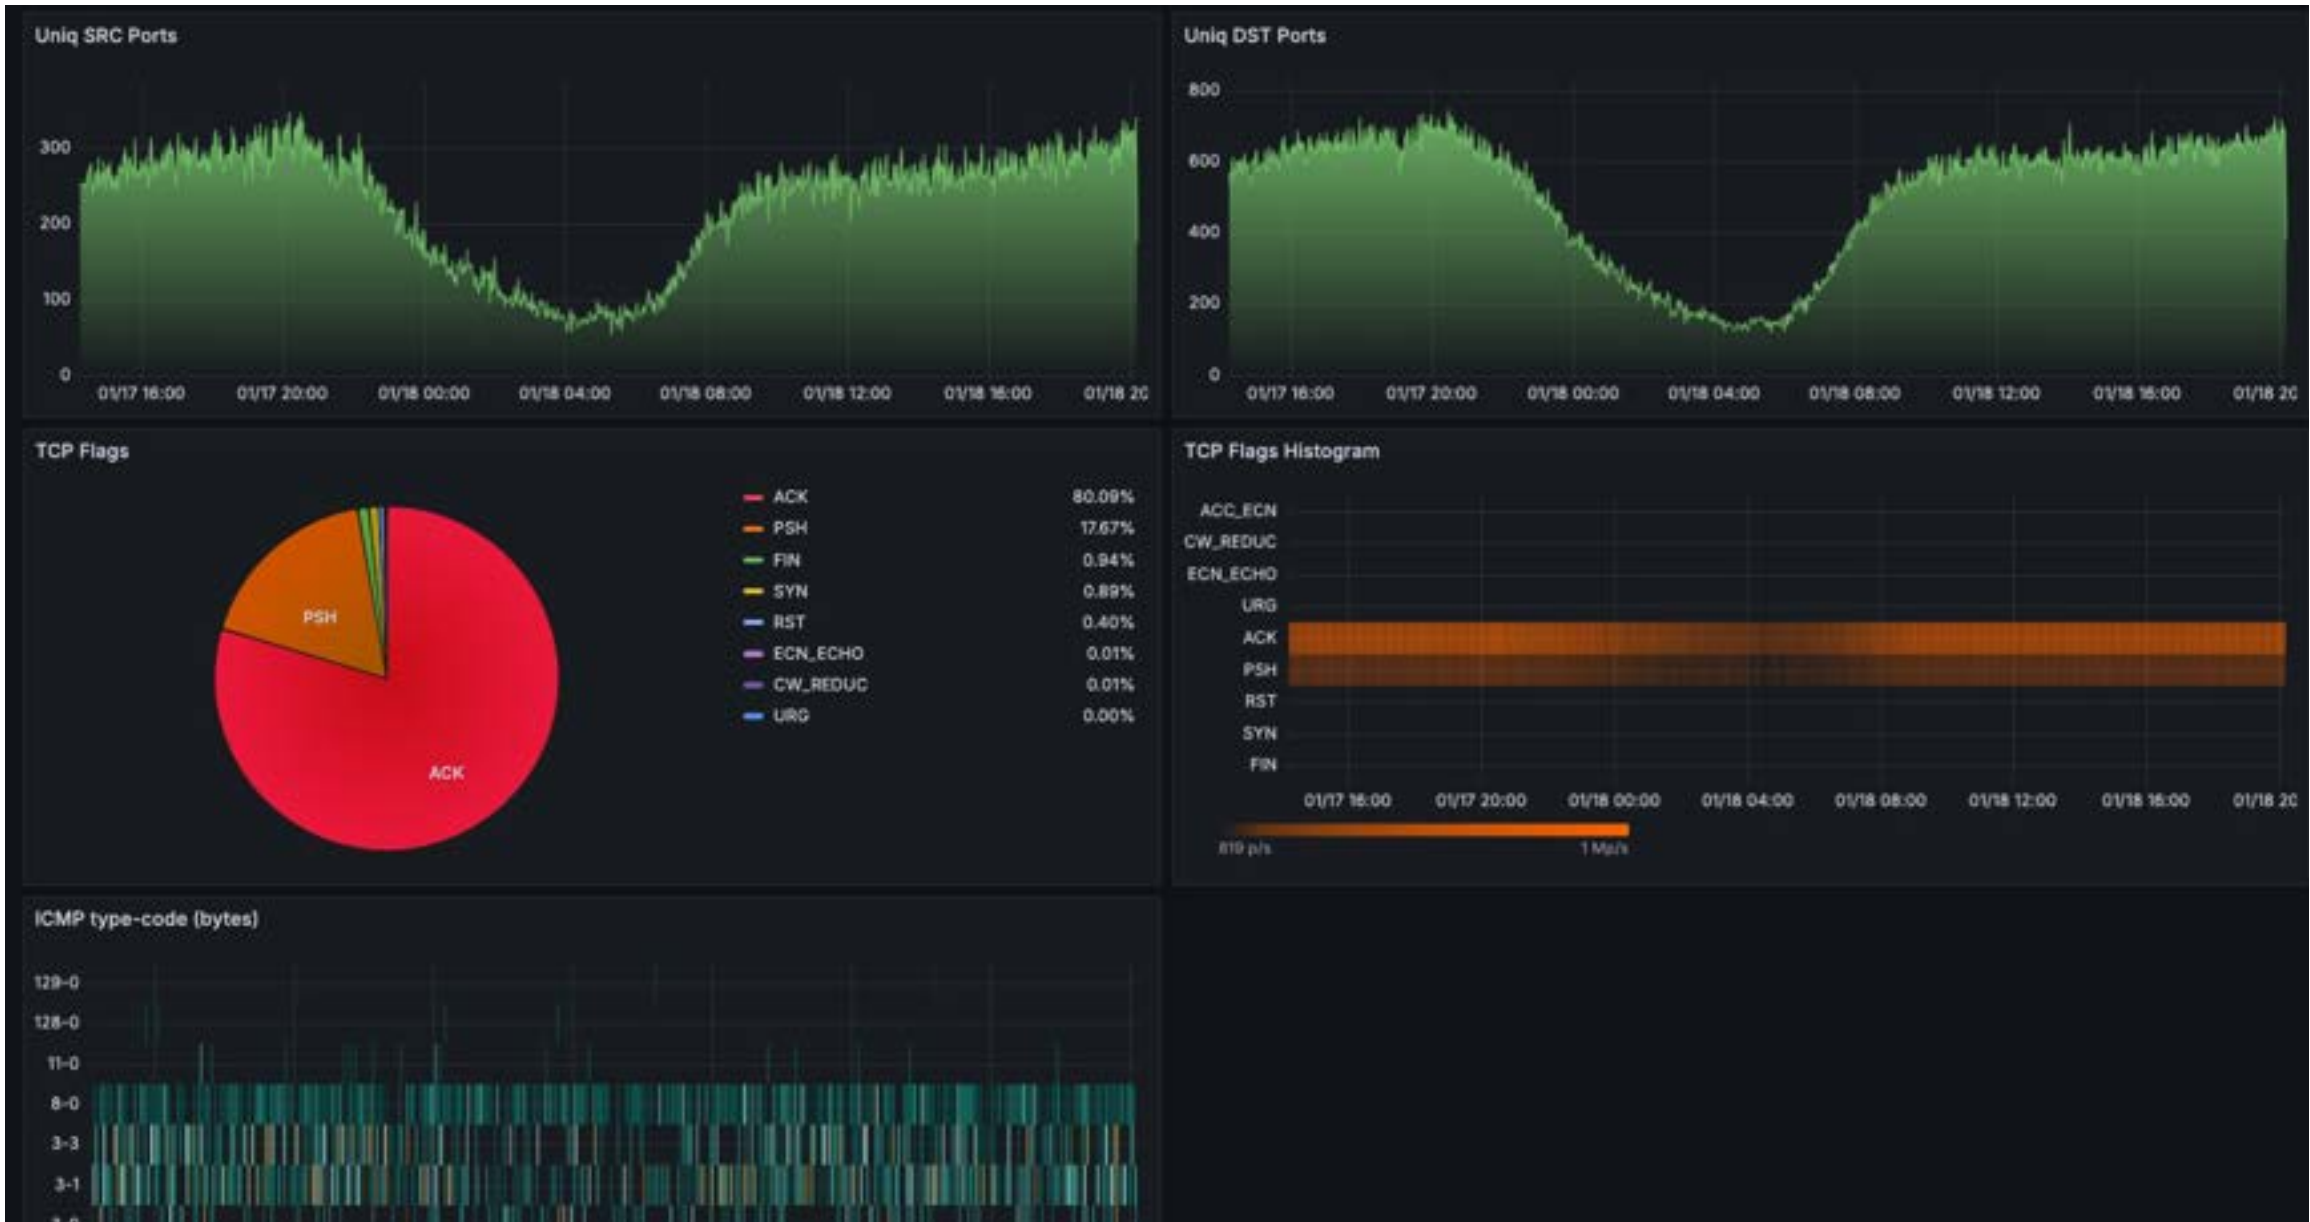


DST Proto+Port Histogram



Uniq SRC Ports

Uniq DST Ports







Top TCP SYN to ACK ratio by SRC IP

217.119.126.72
 217.119.126.250
 212.37.67.11
 195.242.182.247
 185.251.120.196
 185.251.120.157
 185.251.120.118
 185.188.239.206
 176.107.111.66
 176.107.30.81
 176.107.30.100
 176.107.26.7
 176.107.26.236
 176.107.26.181
 172.253.63.94
 172.217.29.67
 172.217.16.195
 142.251.167.94
 142.250.185.206
 142.250.178.3
 142.250.102.108
 91.226.114.197
 91.226.113.151
 91.148.47.12
 87.244.235.129
 87.244.209.8
 87.244.203.8
 87.244.195.182
 62.269.165.206
 46.228.222.25
 5.59.168.18



Traffic flow by MAC address (IXP mainly)

Traffic source by SRC MAC (IXP mainly)



Name	Mean
AS15169 - GOOGLE - IX NIX.CZ (00:19:E2:D1:8F:F0)	874 Mb/s
AS15169 - GOOGLE - IX NIX.SK (EC:13:D8:78:1F:CA)	867 Mb/s
AS15169 - GOOGLE - IX NIX.SK (EC:13:D8:79:89:23)	868 Mb/s
AS15169 - GOOGLE - IX PEERING.CZ (EC:13:D8:79:8F:C8)	866 Mb/s
AS15169 - GOOGLE - IX NIX.CZ (EC:13:D8:79:8F:C3)	853 Mb/s
AS5511 - Orange - TRANZIT ORANGE	9.10 Mb/s
AS39392 - SH.cz s.r.o. - TRANZIT PEERING.CZ	4.94 Mb/s
AS39392 - SH.cz s.r.o. - TRANZIT PEERING.CZ	4.89 Mb/s
AS48326 - DataNetworks s.r. - IX SIX.SK (48:FD:8E:90:FF:71)	3.95 Mb/s
AS6855 - Slovak Telekom, a - IX SIX.SK (F8:39:18:78:E8:17)	1.42 Mb/s

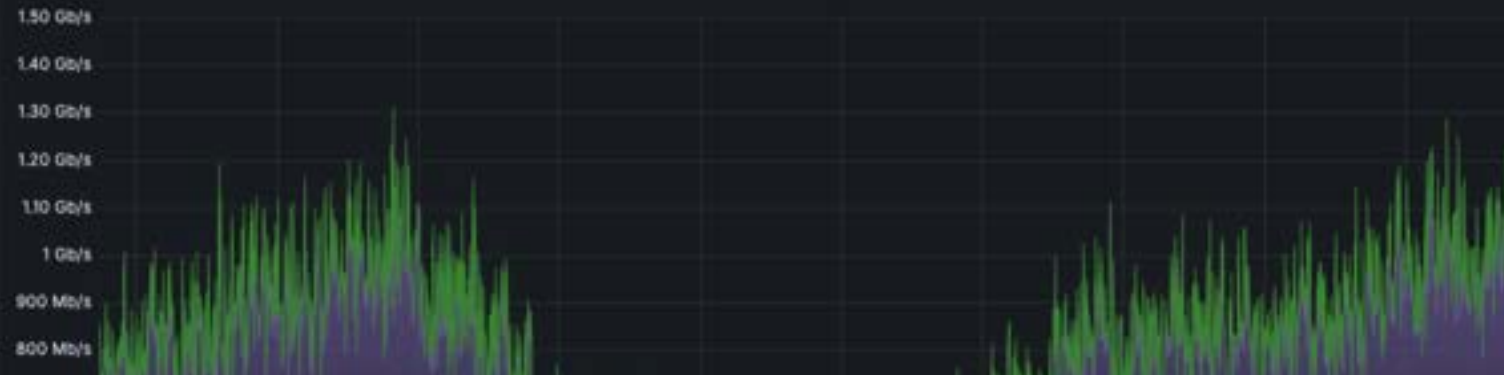
~ Traffic flow by MAC address (IXP mainly)

Traffic source by SRC MAC (IXP mainly)



Name	Mean
AS15169 - GOOGLE - IX NIX.CZ (00:19:E2:D1:8F:F0)	874 Mb/s
AS15169 - GOOGLE - IX NIX.SK (EC:13:D8:78:1F:CA)	867 Mb/s
AS15169 - GOOGLE - IX NIX.SK (EC:13:D8:79:89:23)	668 Mb/s
AS15169 - GOOGLE - IX PEERING.CZ (EC:13:D8:79:8F:C8)	666 Mb/s
AS15169 - GOOGLE - IX NIX.CZ (EC:13:D8:79:8F:C3)	653 Mb/s
AS5511 - Orange - TRANZIT ORANGE (████████████████████)	9.10 Mb/s
AS39392 - SH.cz s.r.o. - TRANZIT PEERING.CZ (████████████████████)	4.94 Mb/s
AS39392 - SH.cz s.r.o. - TRANZIT PEERING.CZ (████████████████████)	4.89 Mb/s
AS48326 - DataNetworks s.r. - IX SIX.SK (48:FD:8E:90:FF:71)	3.95 Mb/s
AS6855 - Slovak Telekom, a - IX SIX.SK (F8:39:18:78:E8:17)	1.42 Mb/s

Traffic destination by DST MAC (IXP mainly)



Name	Mean
AS15169 - GOOGLE - IX NIX.CZ (EC:13:D8:79:8F:C3)	686 Mb/s
AS15169 - GOOGLE - IX PEERING.CZ (EC:13:D8:79:8F:C8)	14.8 Mb/s
AS5511 - Orange - TRANZIT ORANGE (████████████████████)	13.2 Mb/s
AS31117 - ENERHOTEL a.s. - TRANZIT ORANGE (████████████████████)	8.81 Mb/s
AS48326 - DataNetworks s.r. - IX PEERING.CZ (84:9E:F3:AA:01:FF)	3.85 Mb/s
AS5578 - SWAN, a.s. - TRANZIT SWAN (████████████████████)	3.26 Mb/s
AS39392 - SH.cz s.r.o. - TRANZIT PEERING.CZ (████████████████████)	2.53 Mb/s
AS24115 - ? - IX 1-IX-EU Equinix WAW (4C:73:4F:2B:9B:C0)	233 b/s
AS6939 - HURRICANE - IX NIX.CZ (40:88:2F:BC:28:DD)	160 b/s

Traffic destination by DST MAC (IXP mainly)



Name	Mean
AS15169 - GOOGLE - IX NIX.CZ (EC:13:DB:79:8F:C3)	686 Mb/s
AS15169 - GOOGLE - IX PEERING.CZ (EC:13:DB:79:8F:C8)	14.8 Mb/s
AS5511 - Orange - TRANZIT ORANGE [REDACTED]	13.2 Mb/s
AS31117 - ENERGETEL a.s. - TRANZIT ORANGE [REDACTED]	8.81 Mb/s
AS48326 - DataNetworks s.r. - IX PEERING.CZ (64:9E:F3:AA:01:FF)	3.85 Mb/s
AS5578 - SWAN, a.s. - TRANZIT SWAN [REDACTED]	3.26 Mb/s
AS39392 - SH.cz s.r.o. - TRANZIT PEERING.CZ [REDACTED]	2.53 Mb/s
AS24115 - 7 - IX 1-IX-EU Equinix WAW (4C:73:4F:2B:9B:C0)	233 b/s
AS6939 - HURRICANE - IX NIX.CZ (40:88:2F:BC:28:DD)	160 b/s

Sankey AS traffic

Sankey TOP WORLD->ENERGETEL



Sankey TOP ENERGETEL-> WORLD



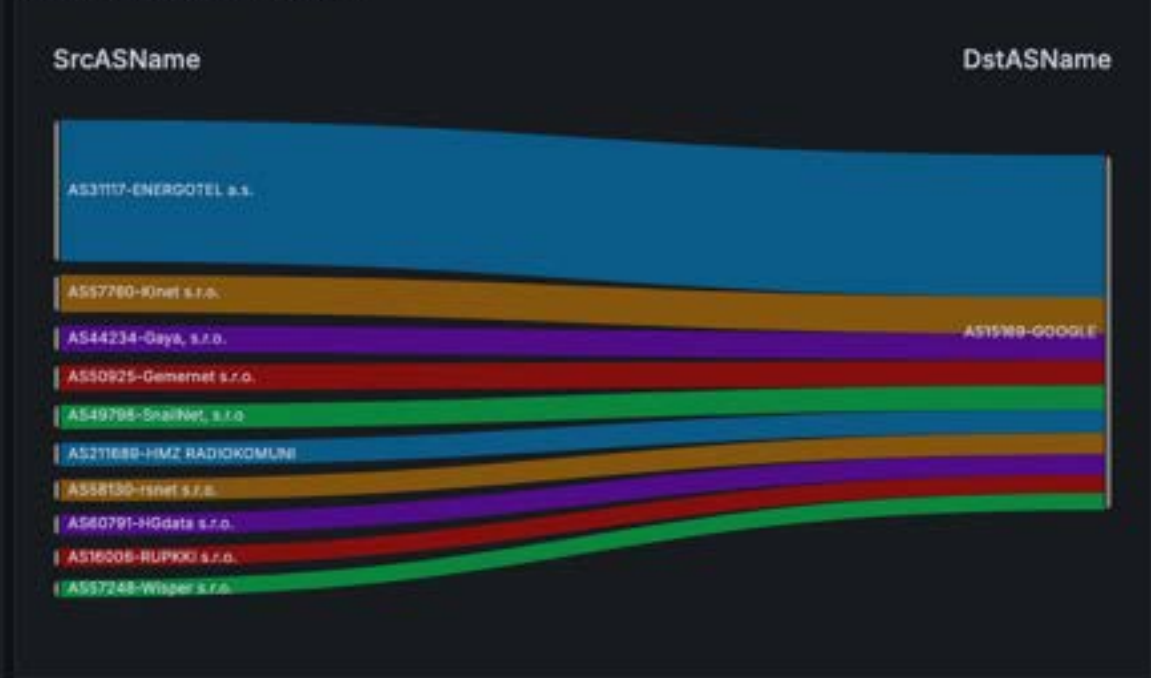


Sankey AS traffic

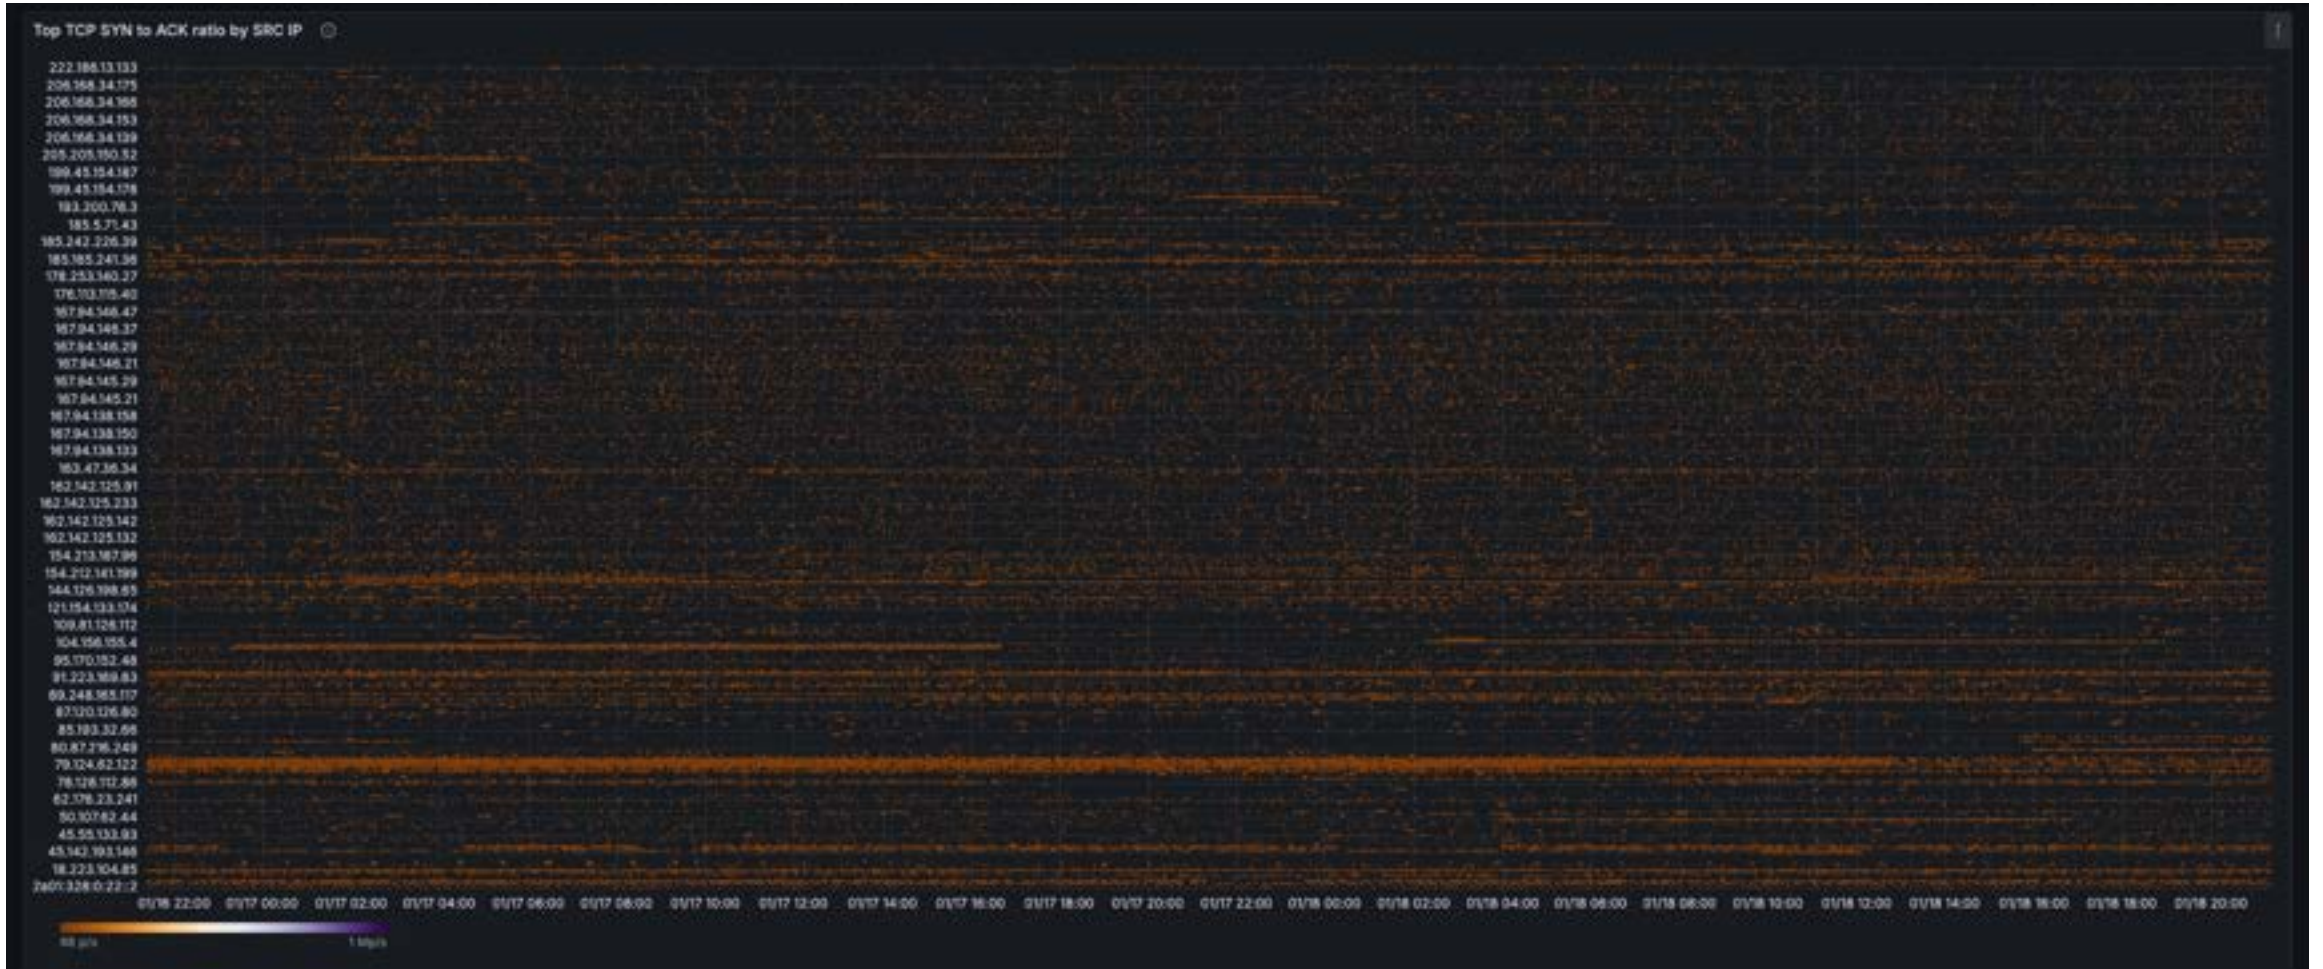
Sankey TOP WORLD->ENERGOTEL



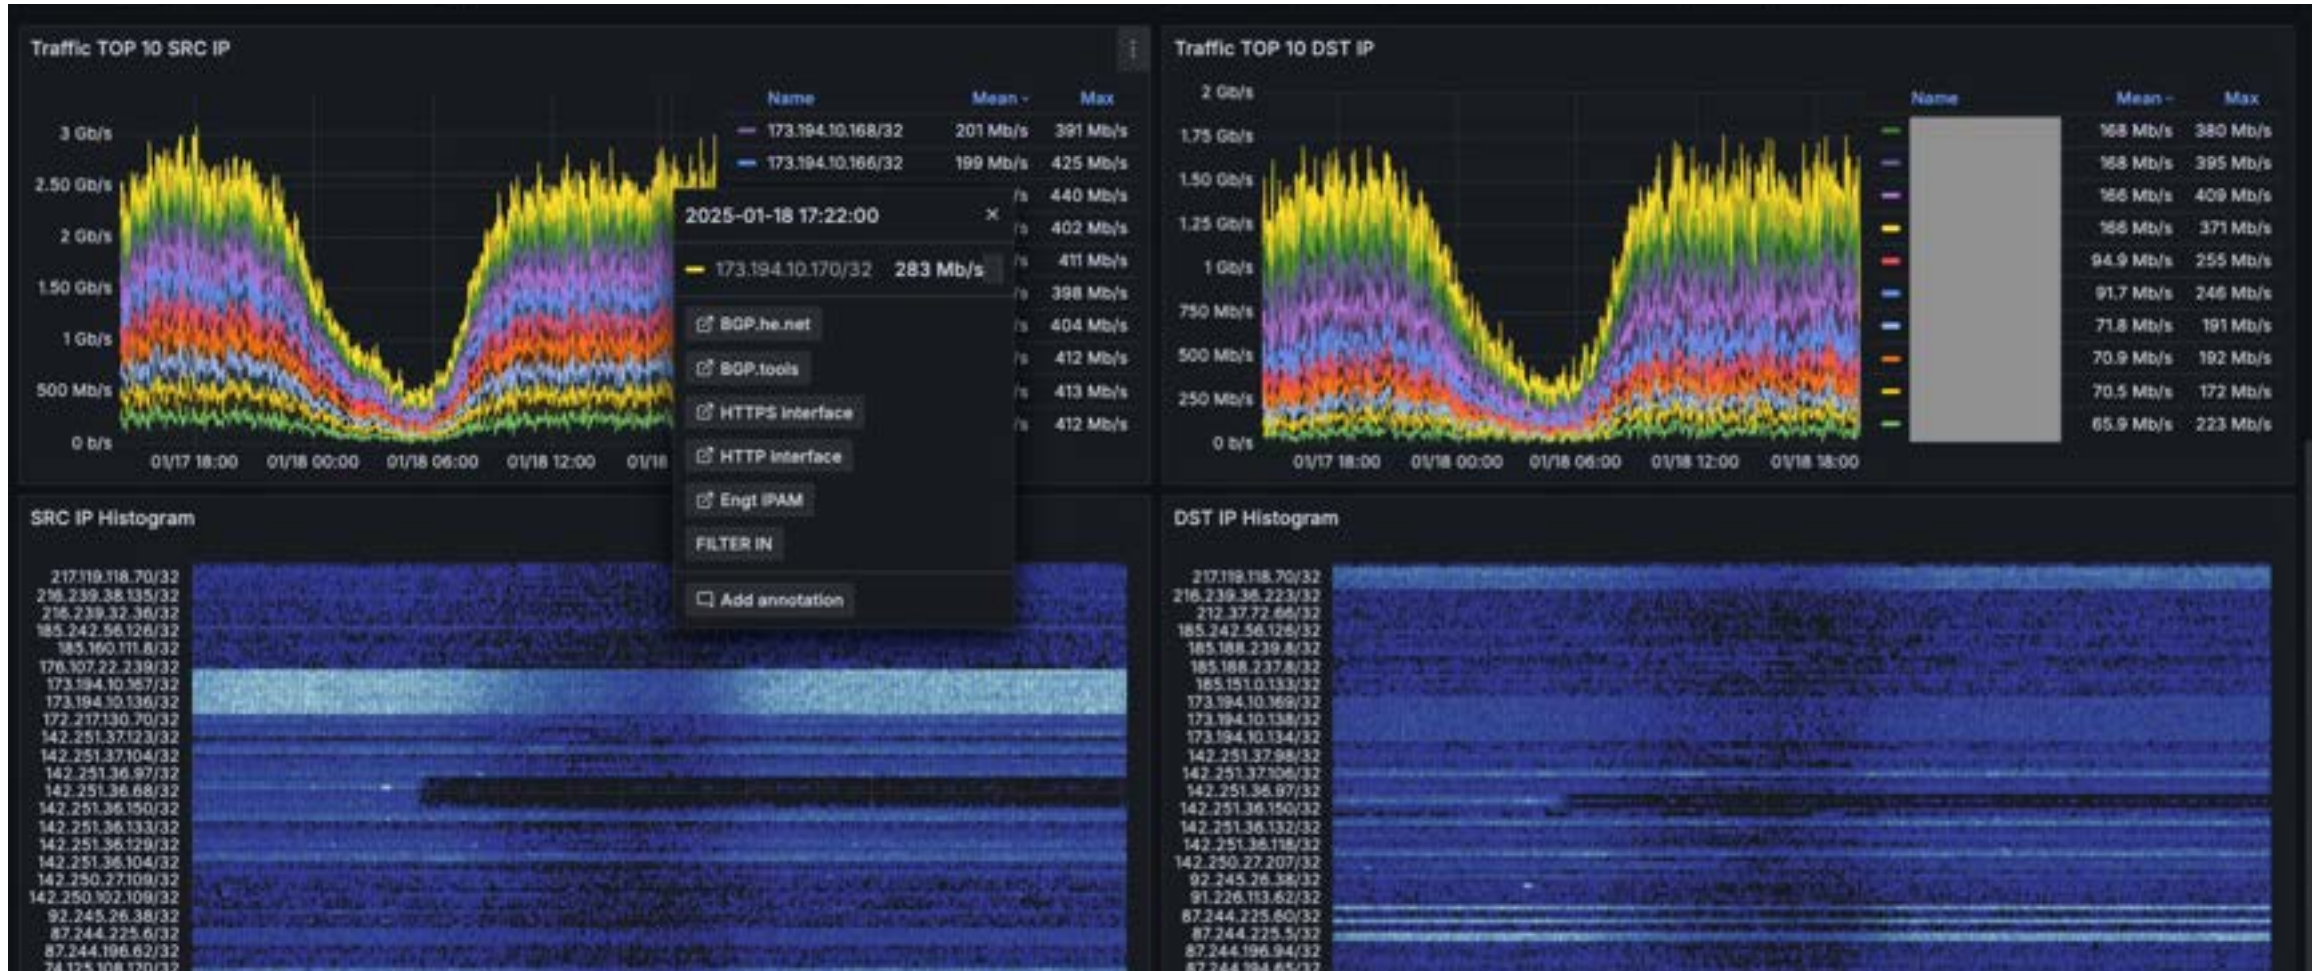
Sankey TOP ENERGOTEL-> WORLD



Port scanners detection

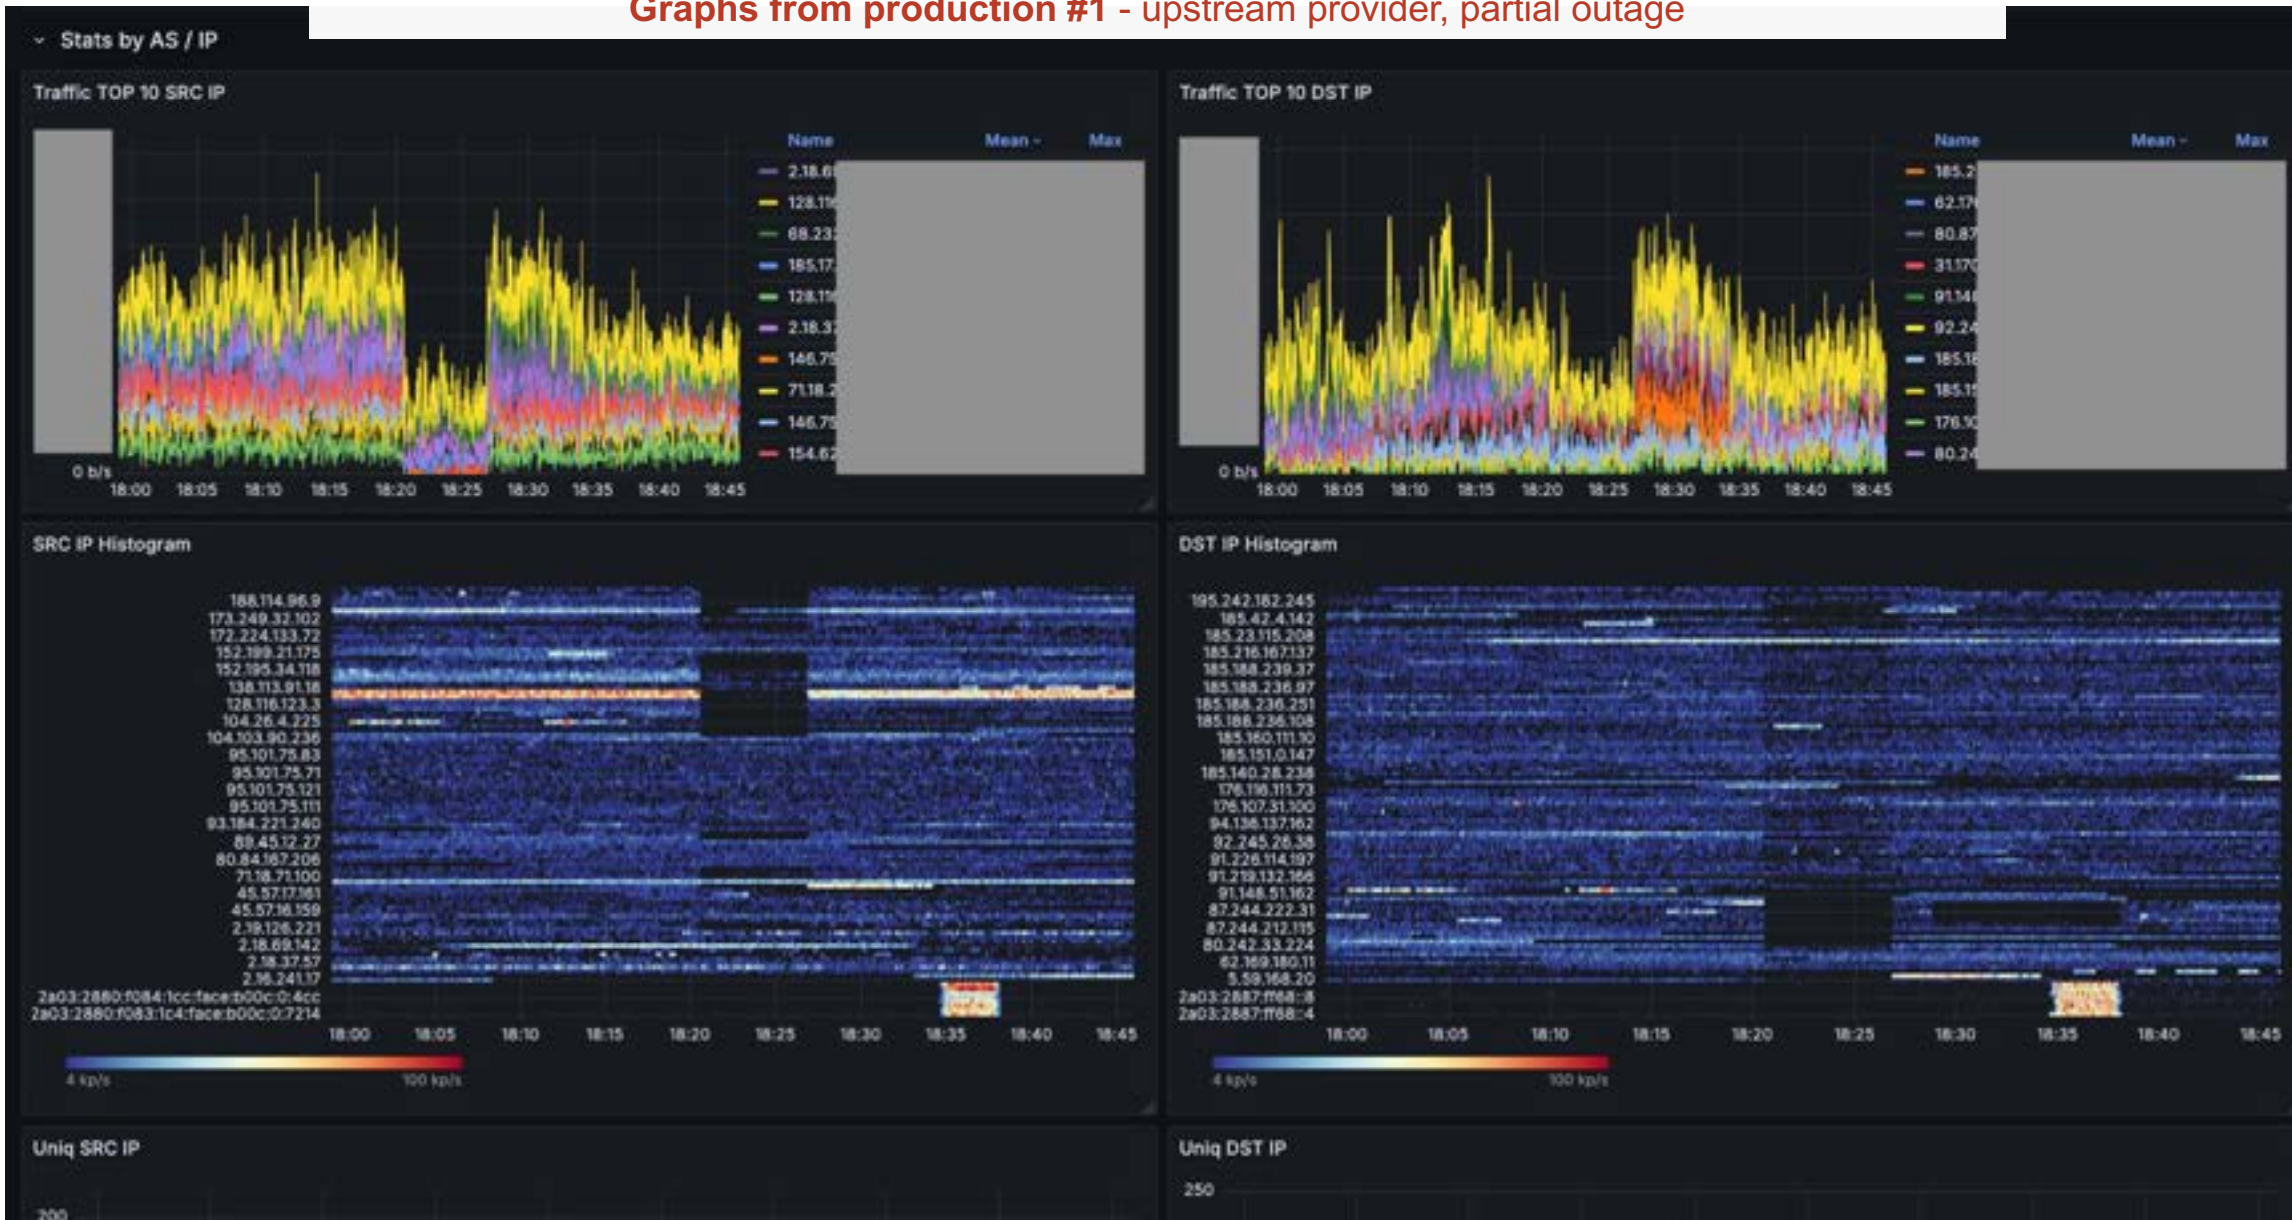


Data links

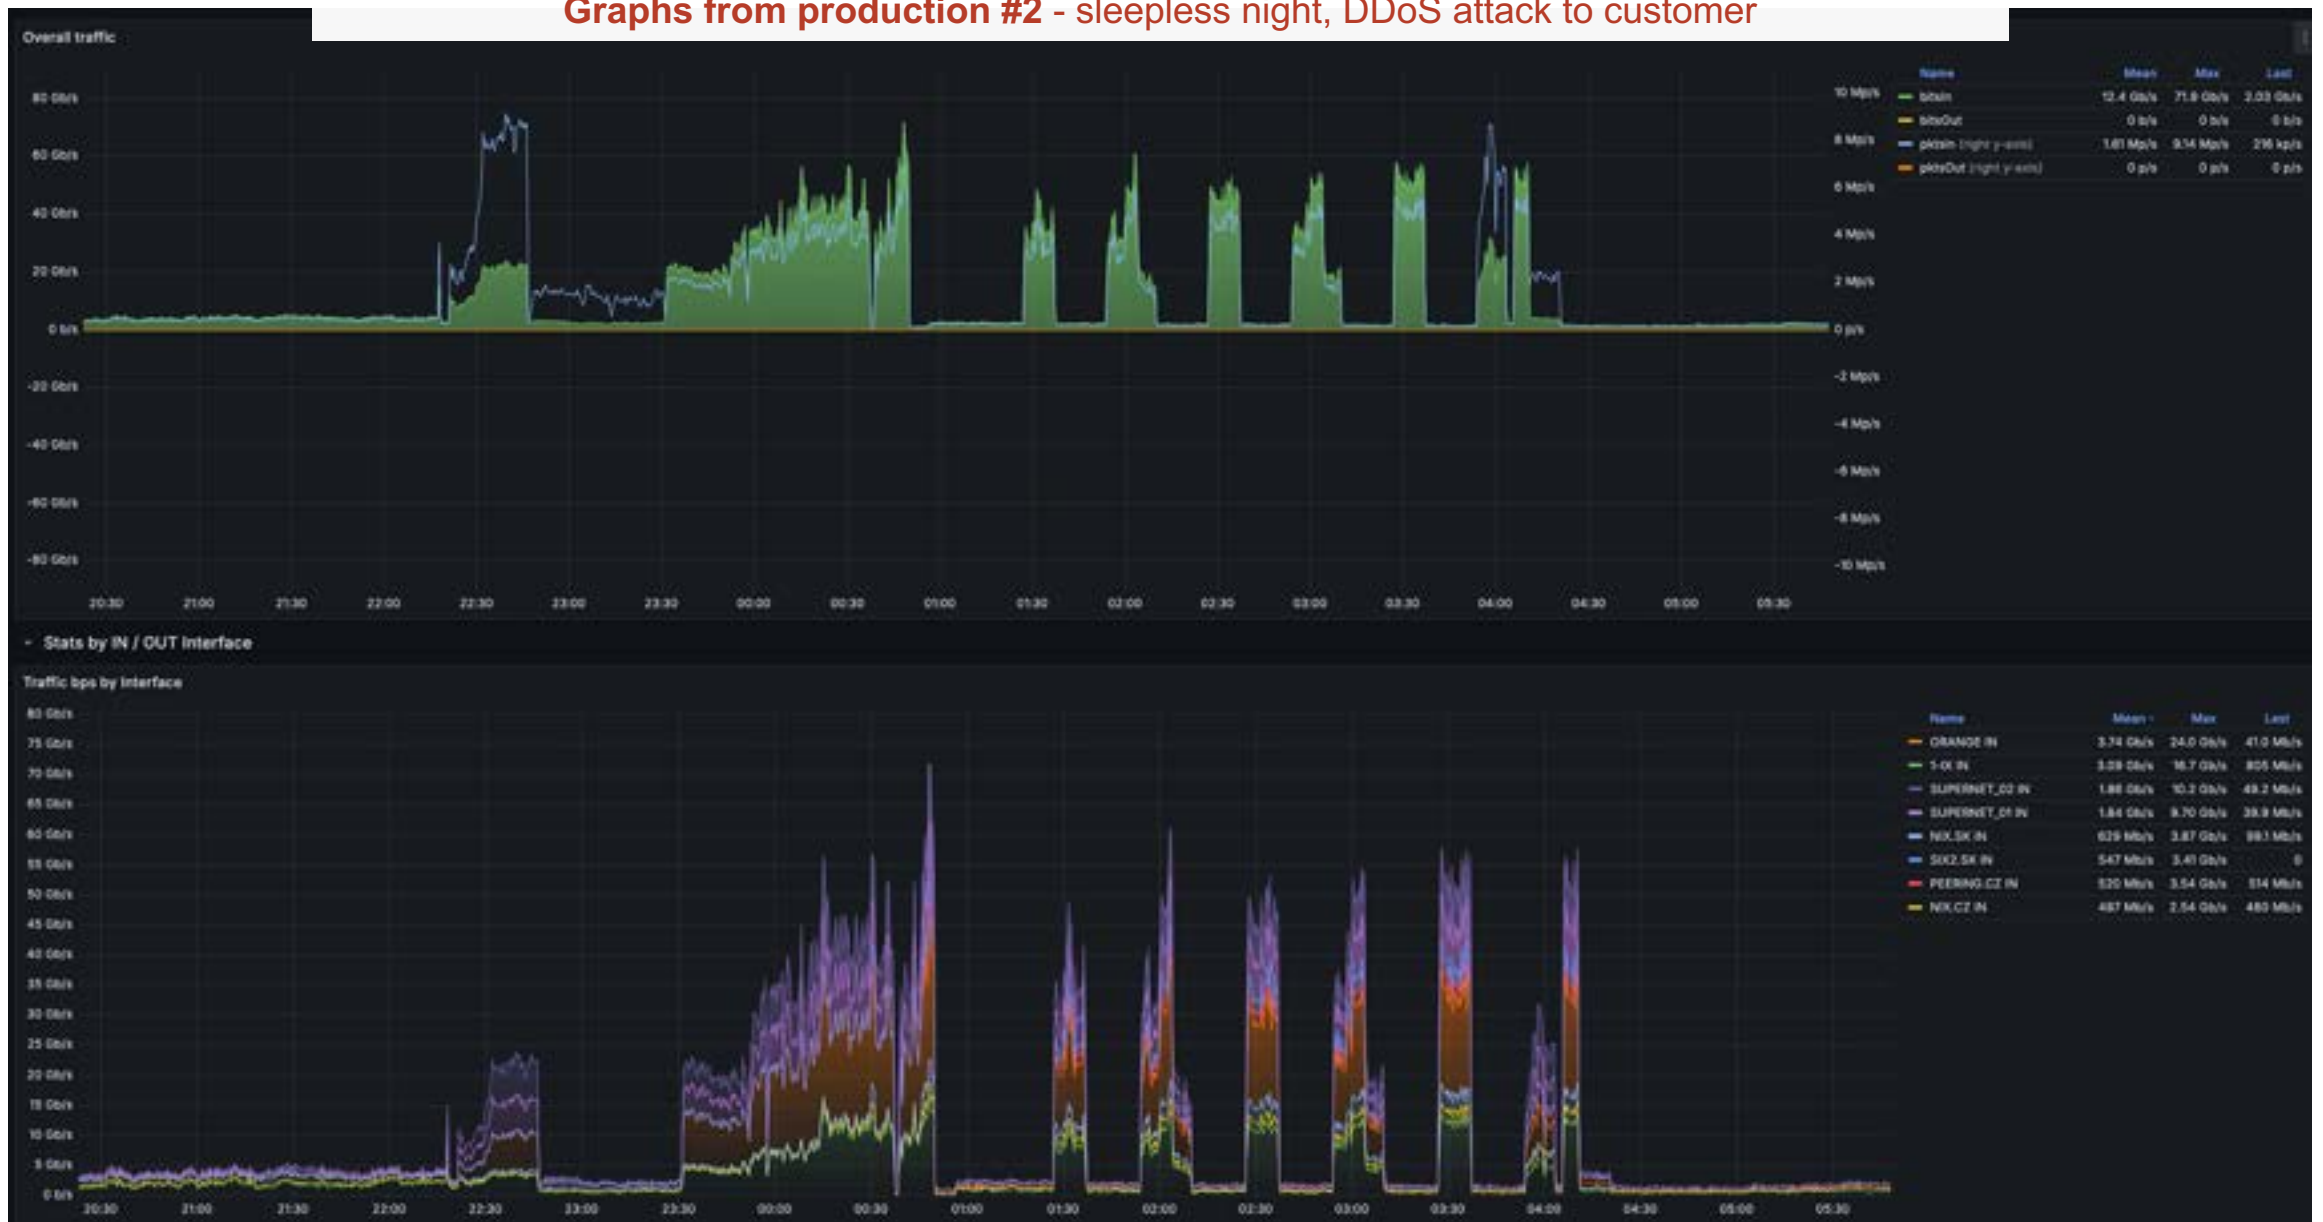


Graphs from production

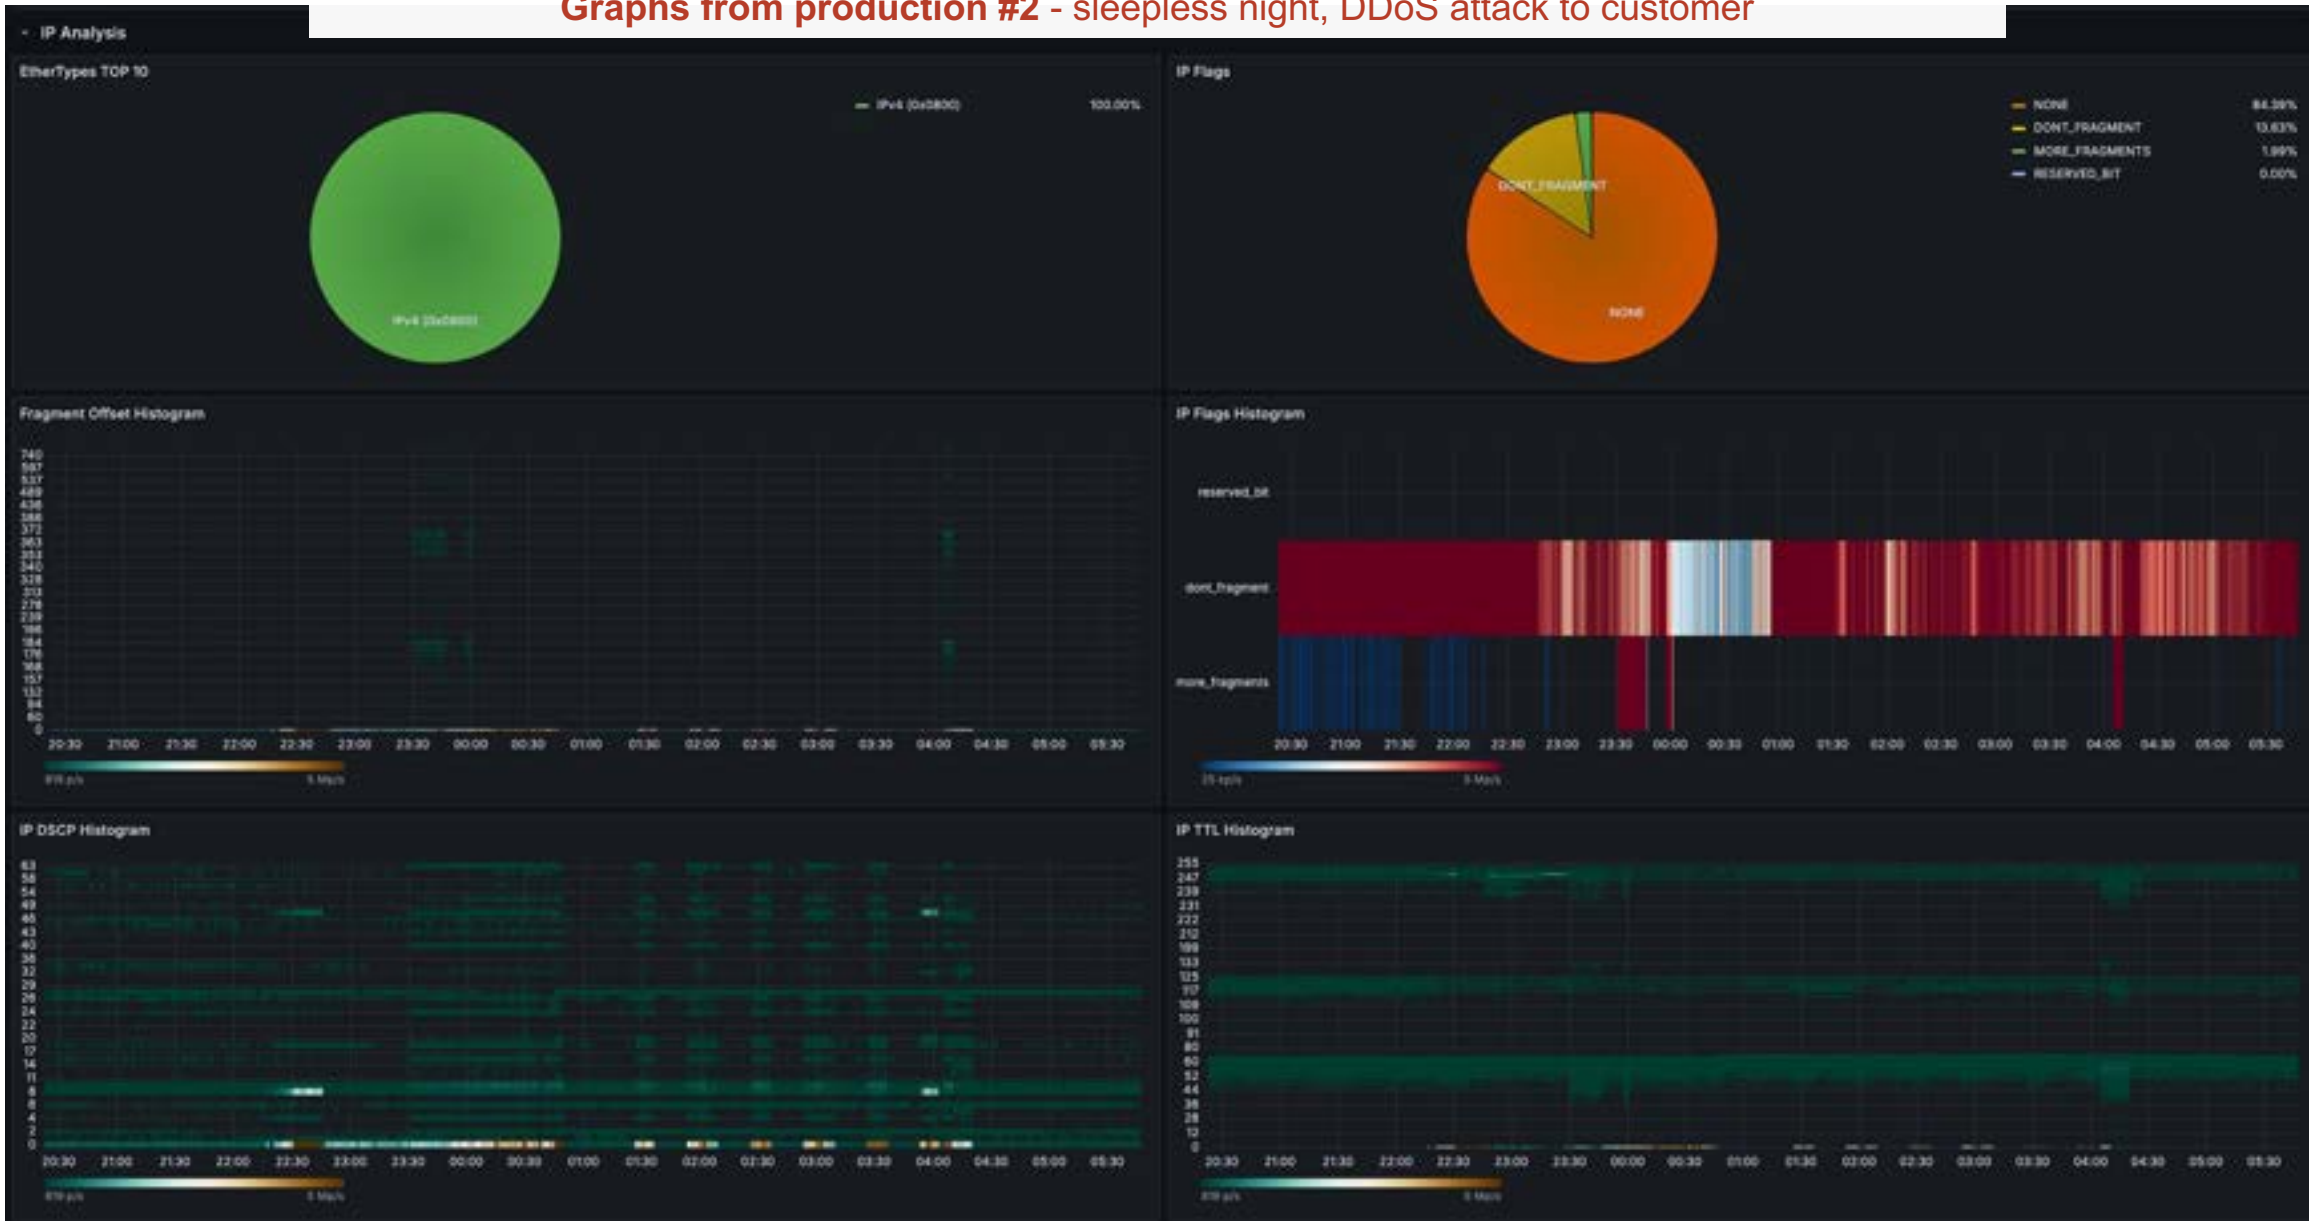
Graphs from production #1 - upstream provider, partial outage



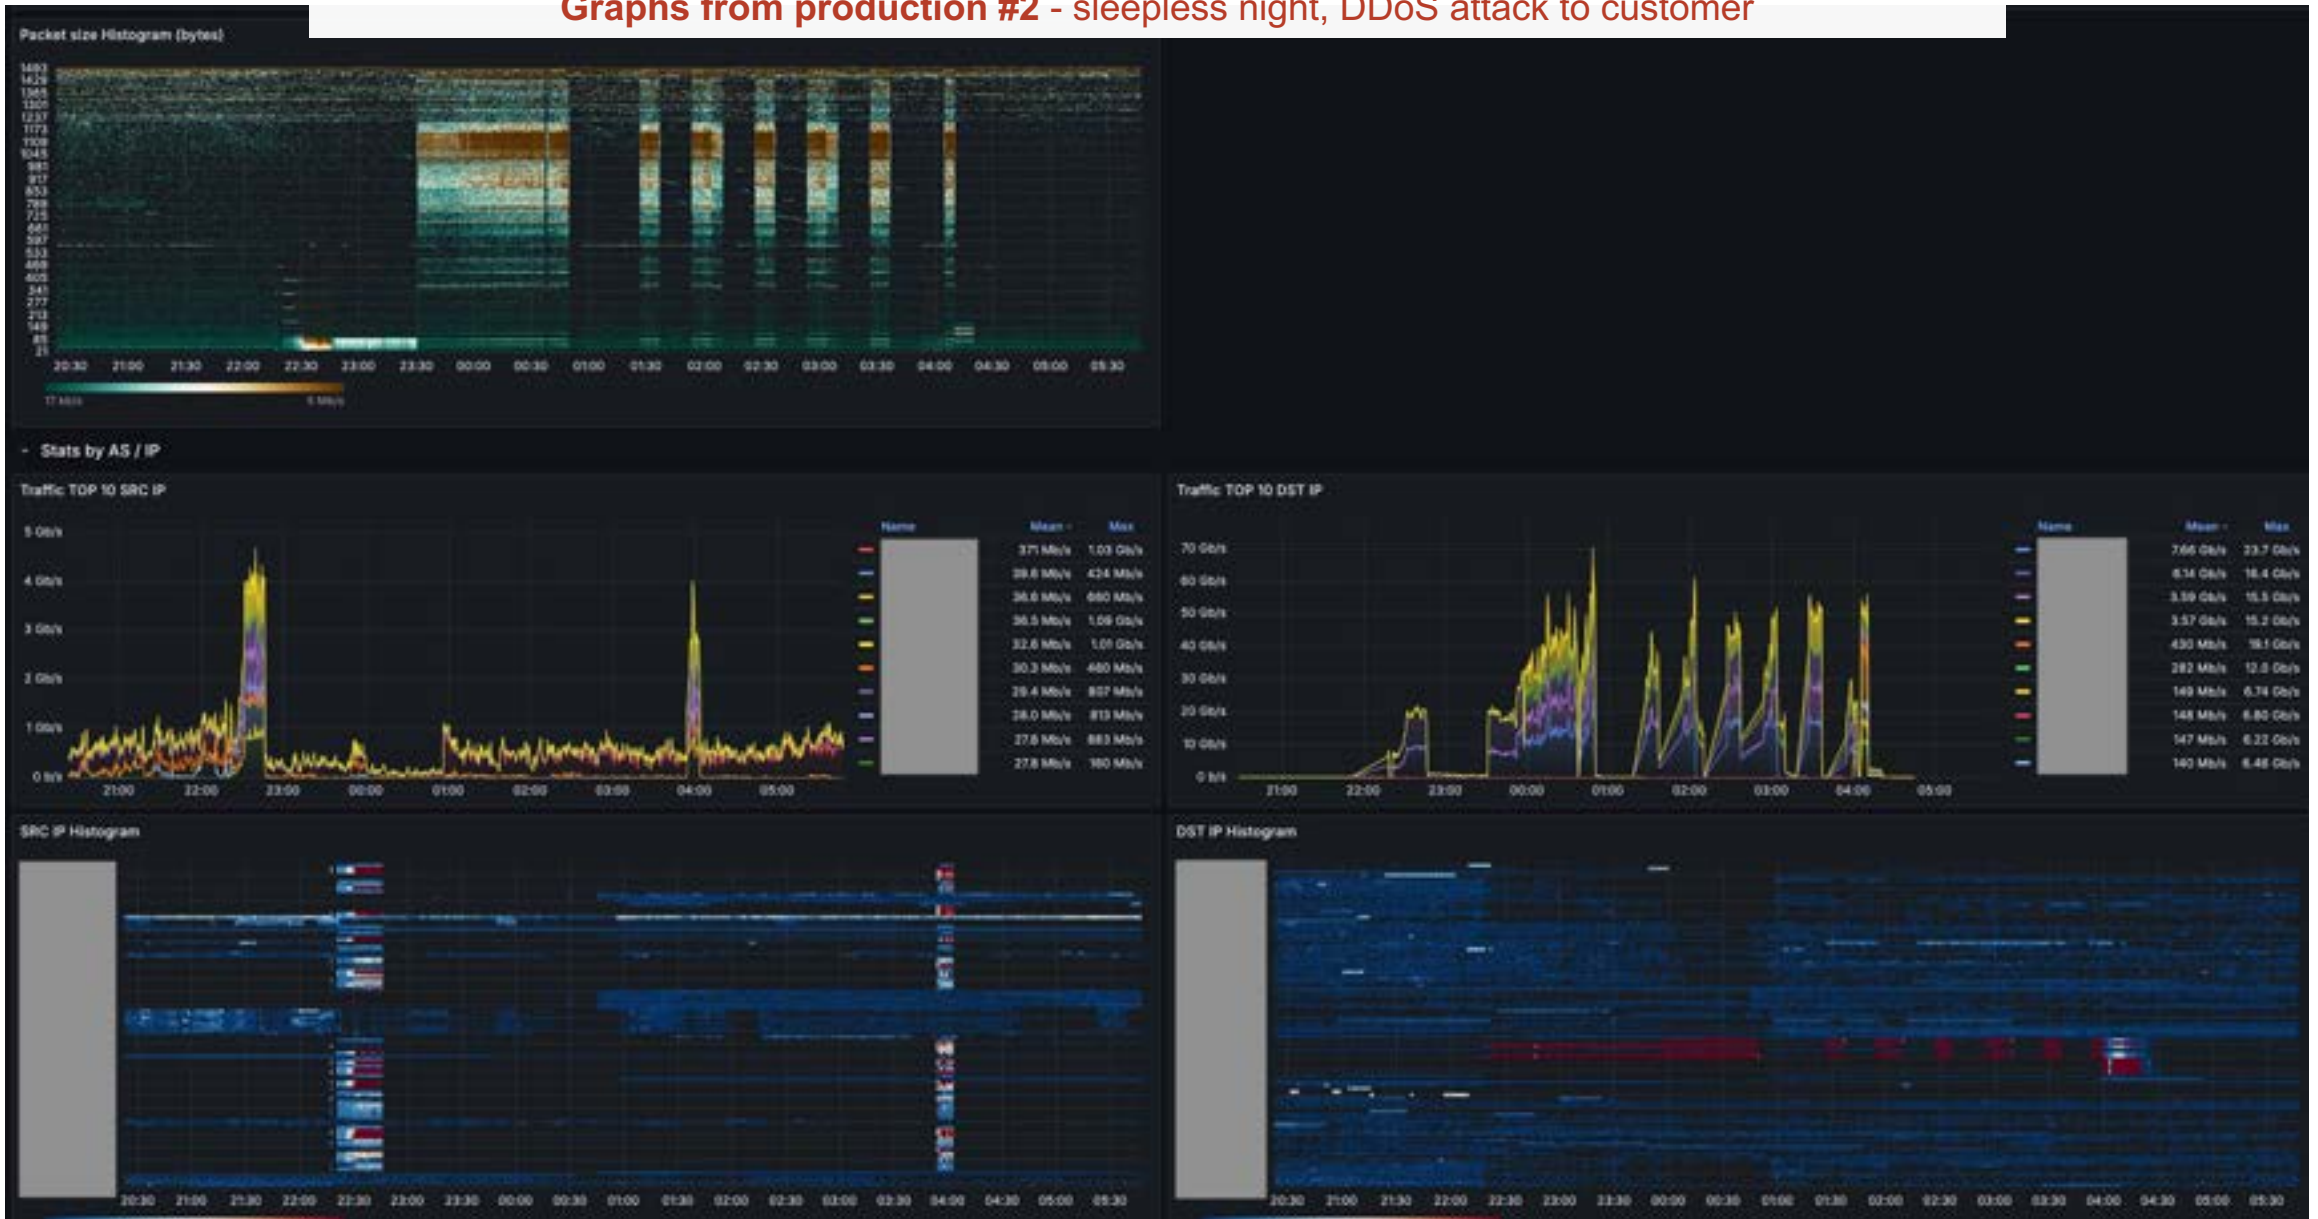
Graphs from production #2 - sleepless night, DDoS attack to customer



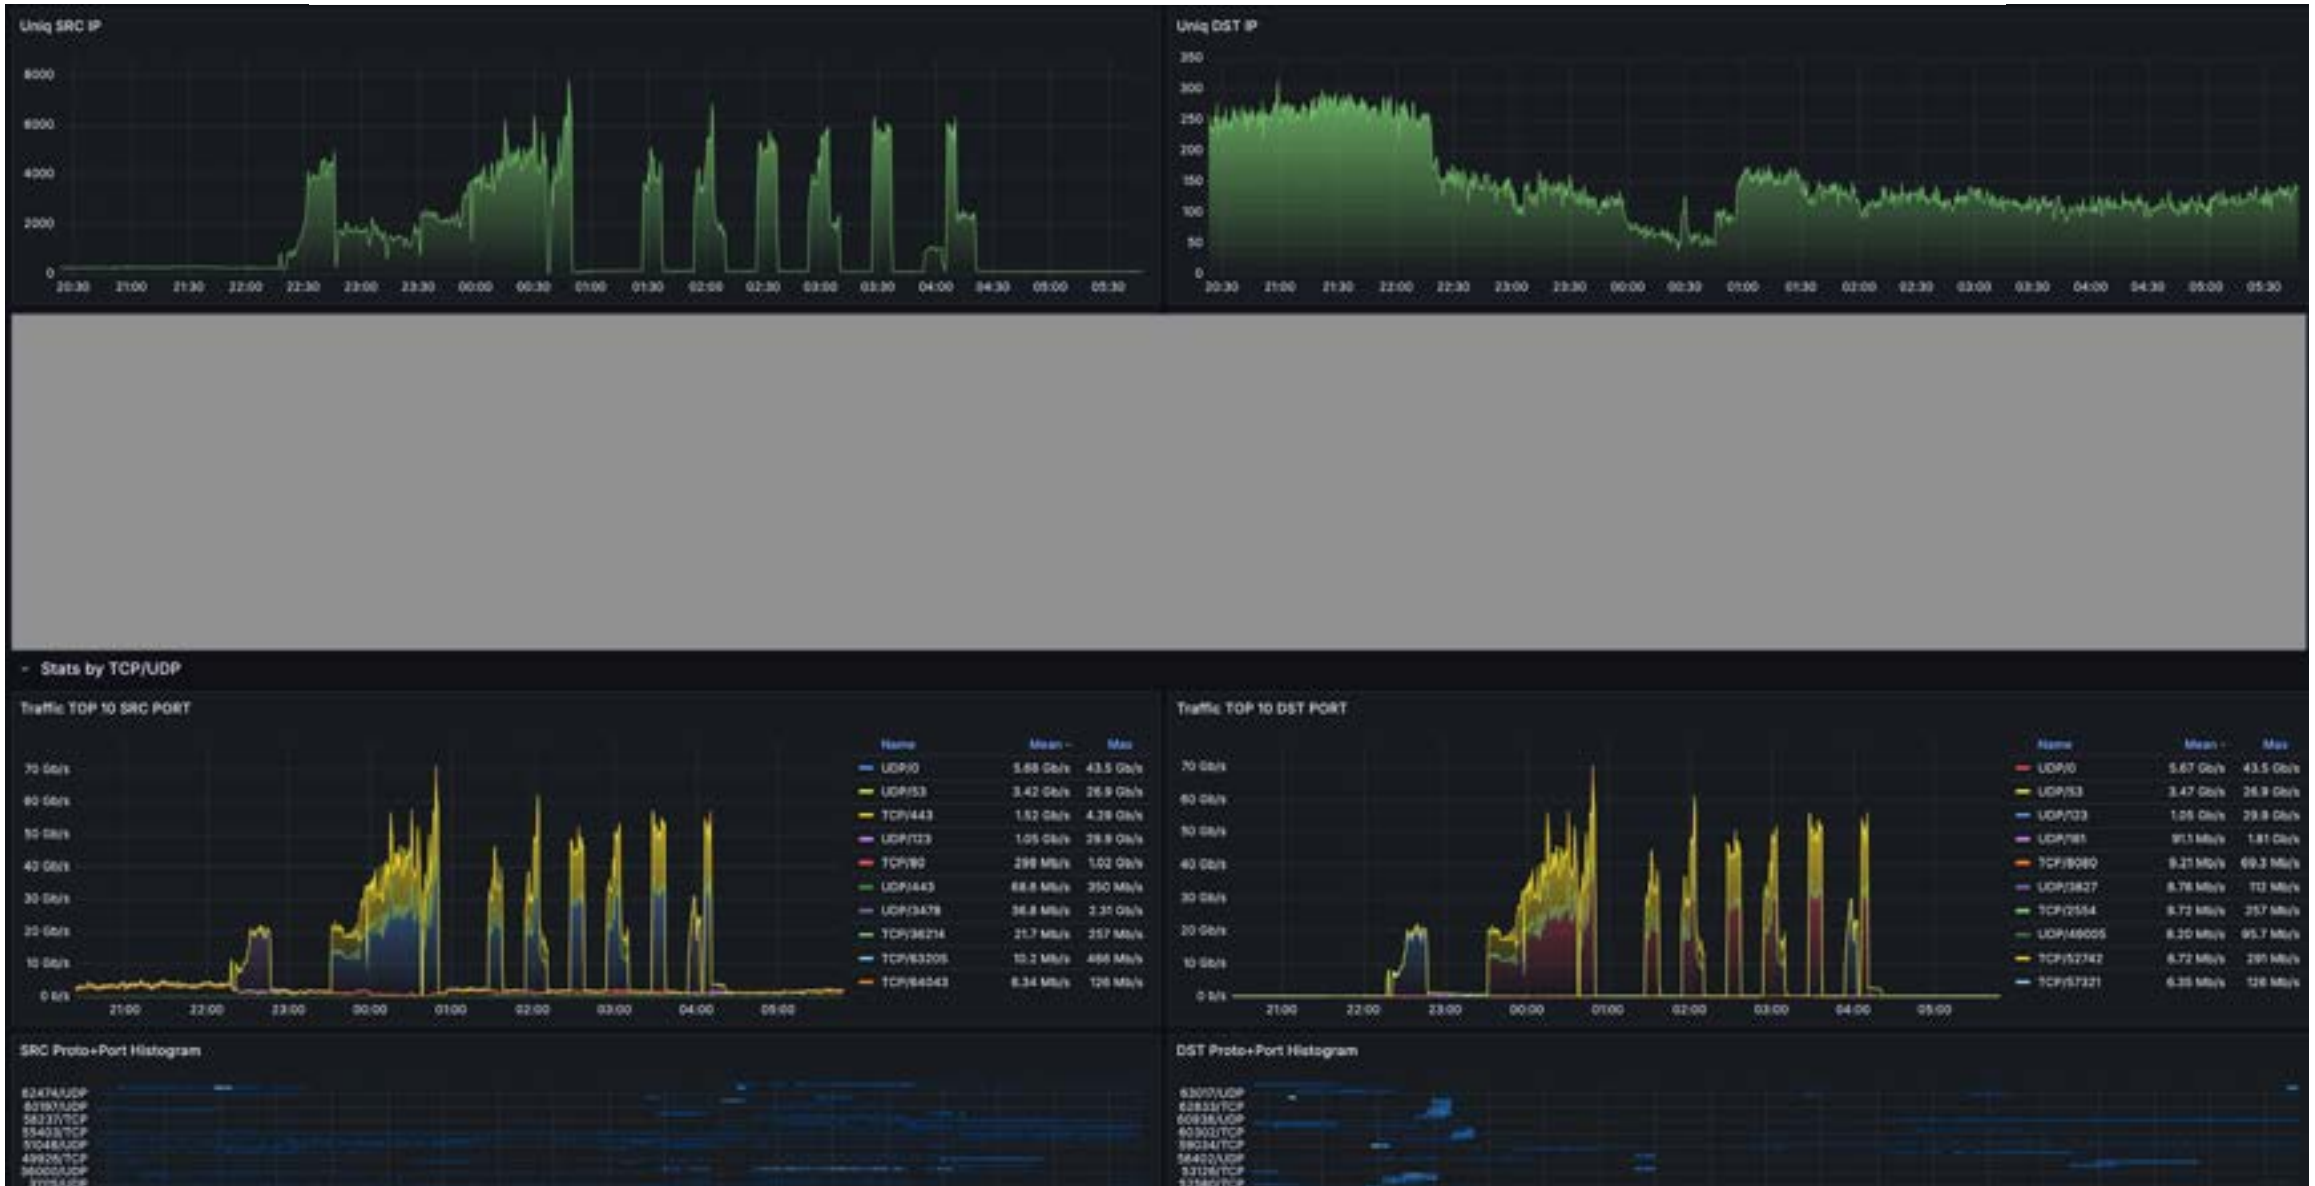
Graphs from production #2 - sleepless night, DDoS attack to customer



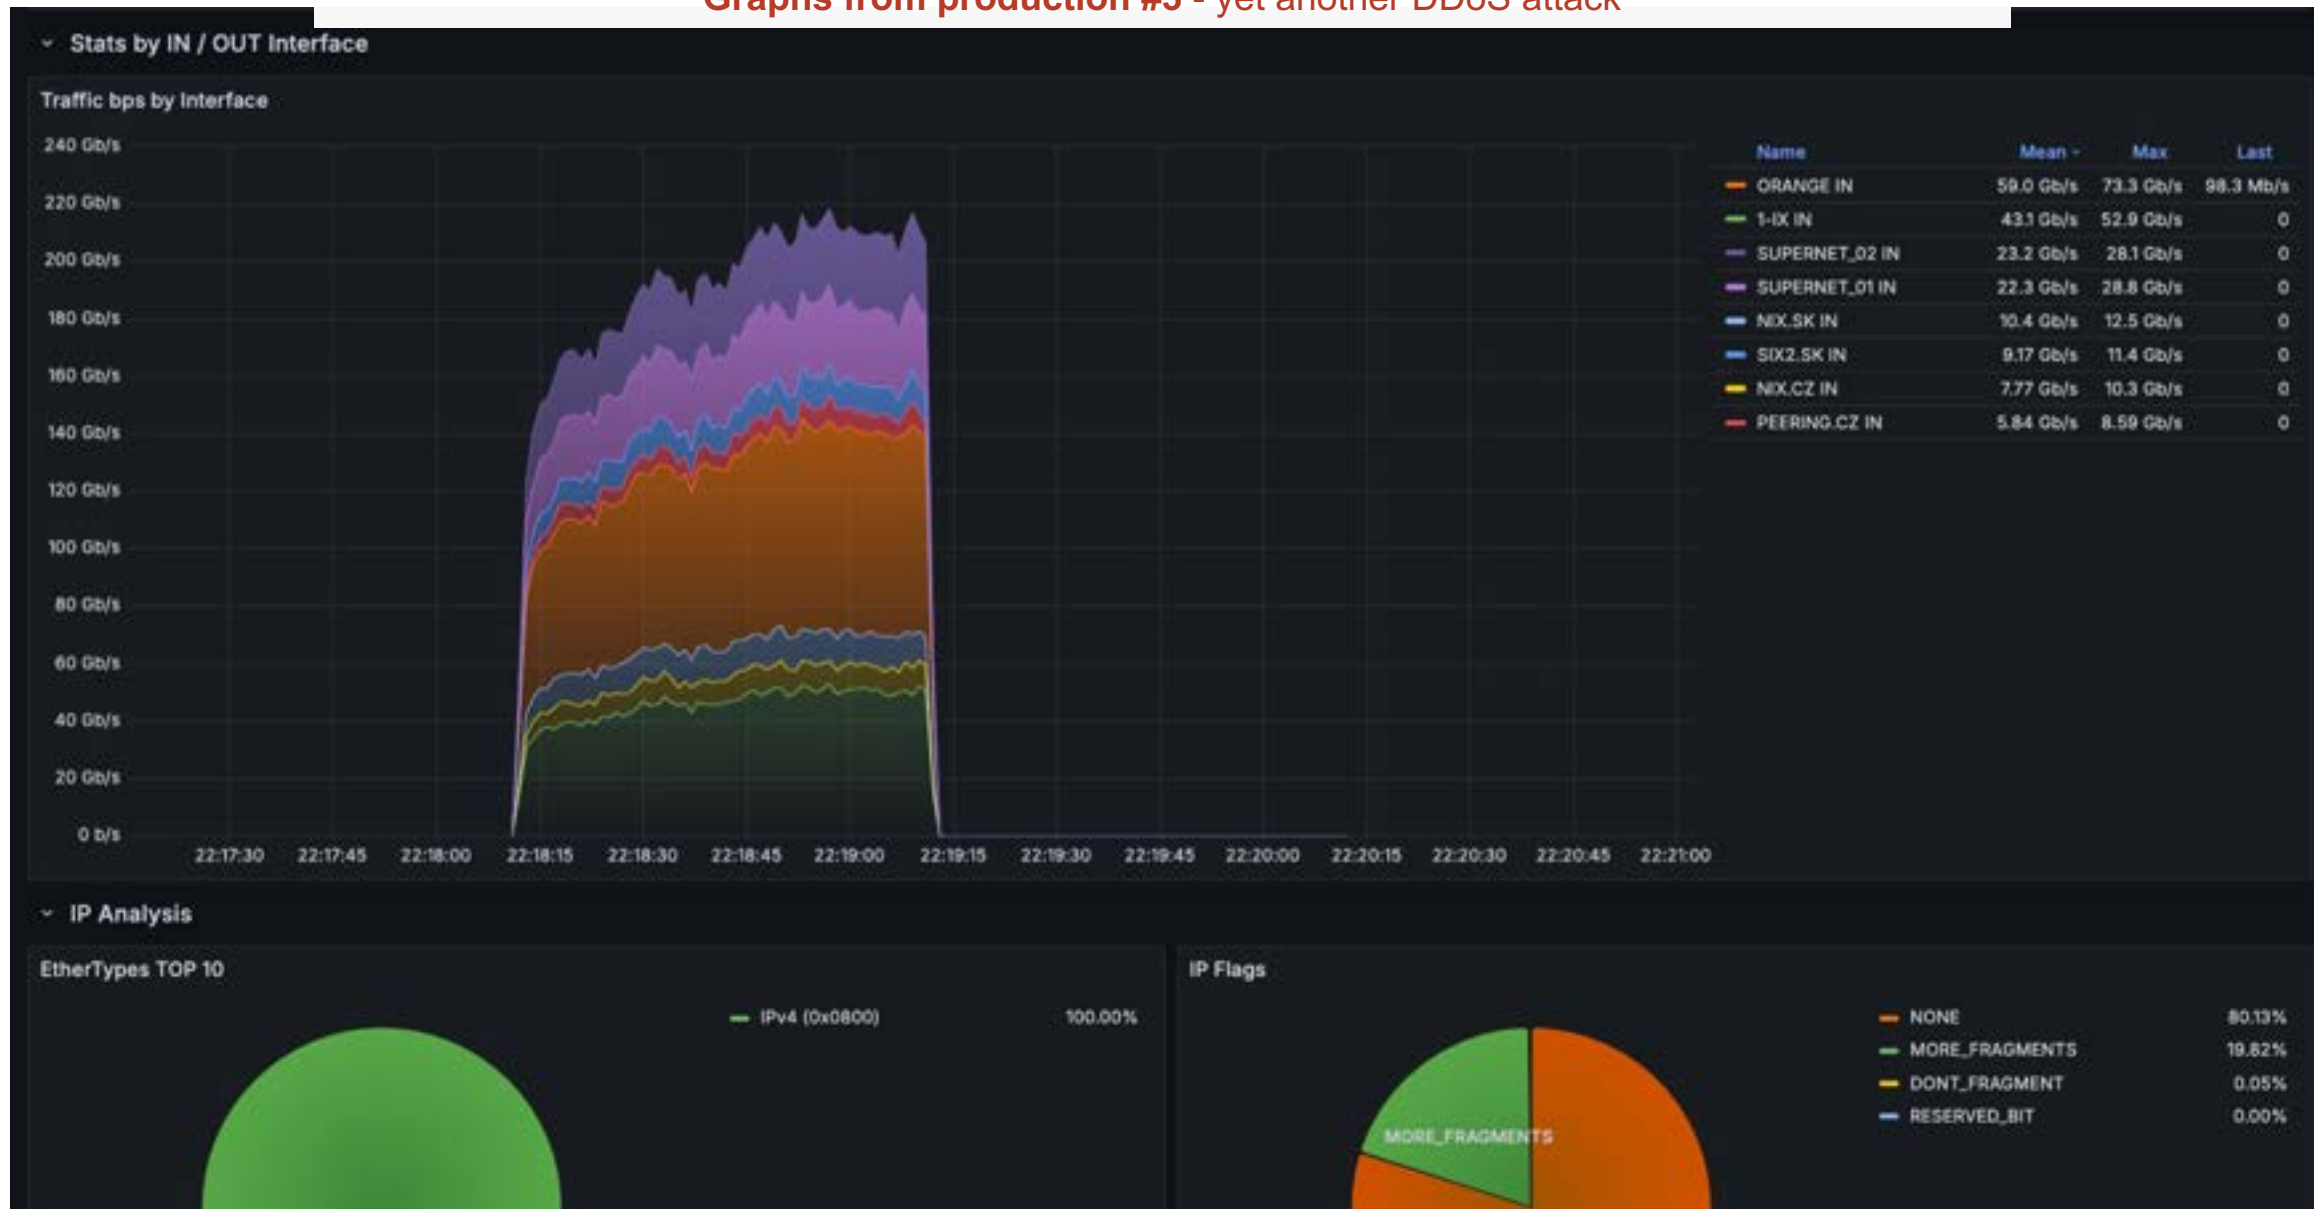
Graphs from production #2 - sleepless night, DDoS attack to customer



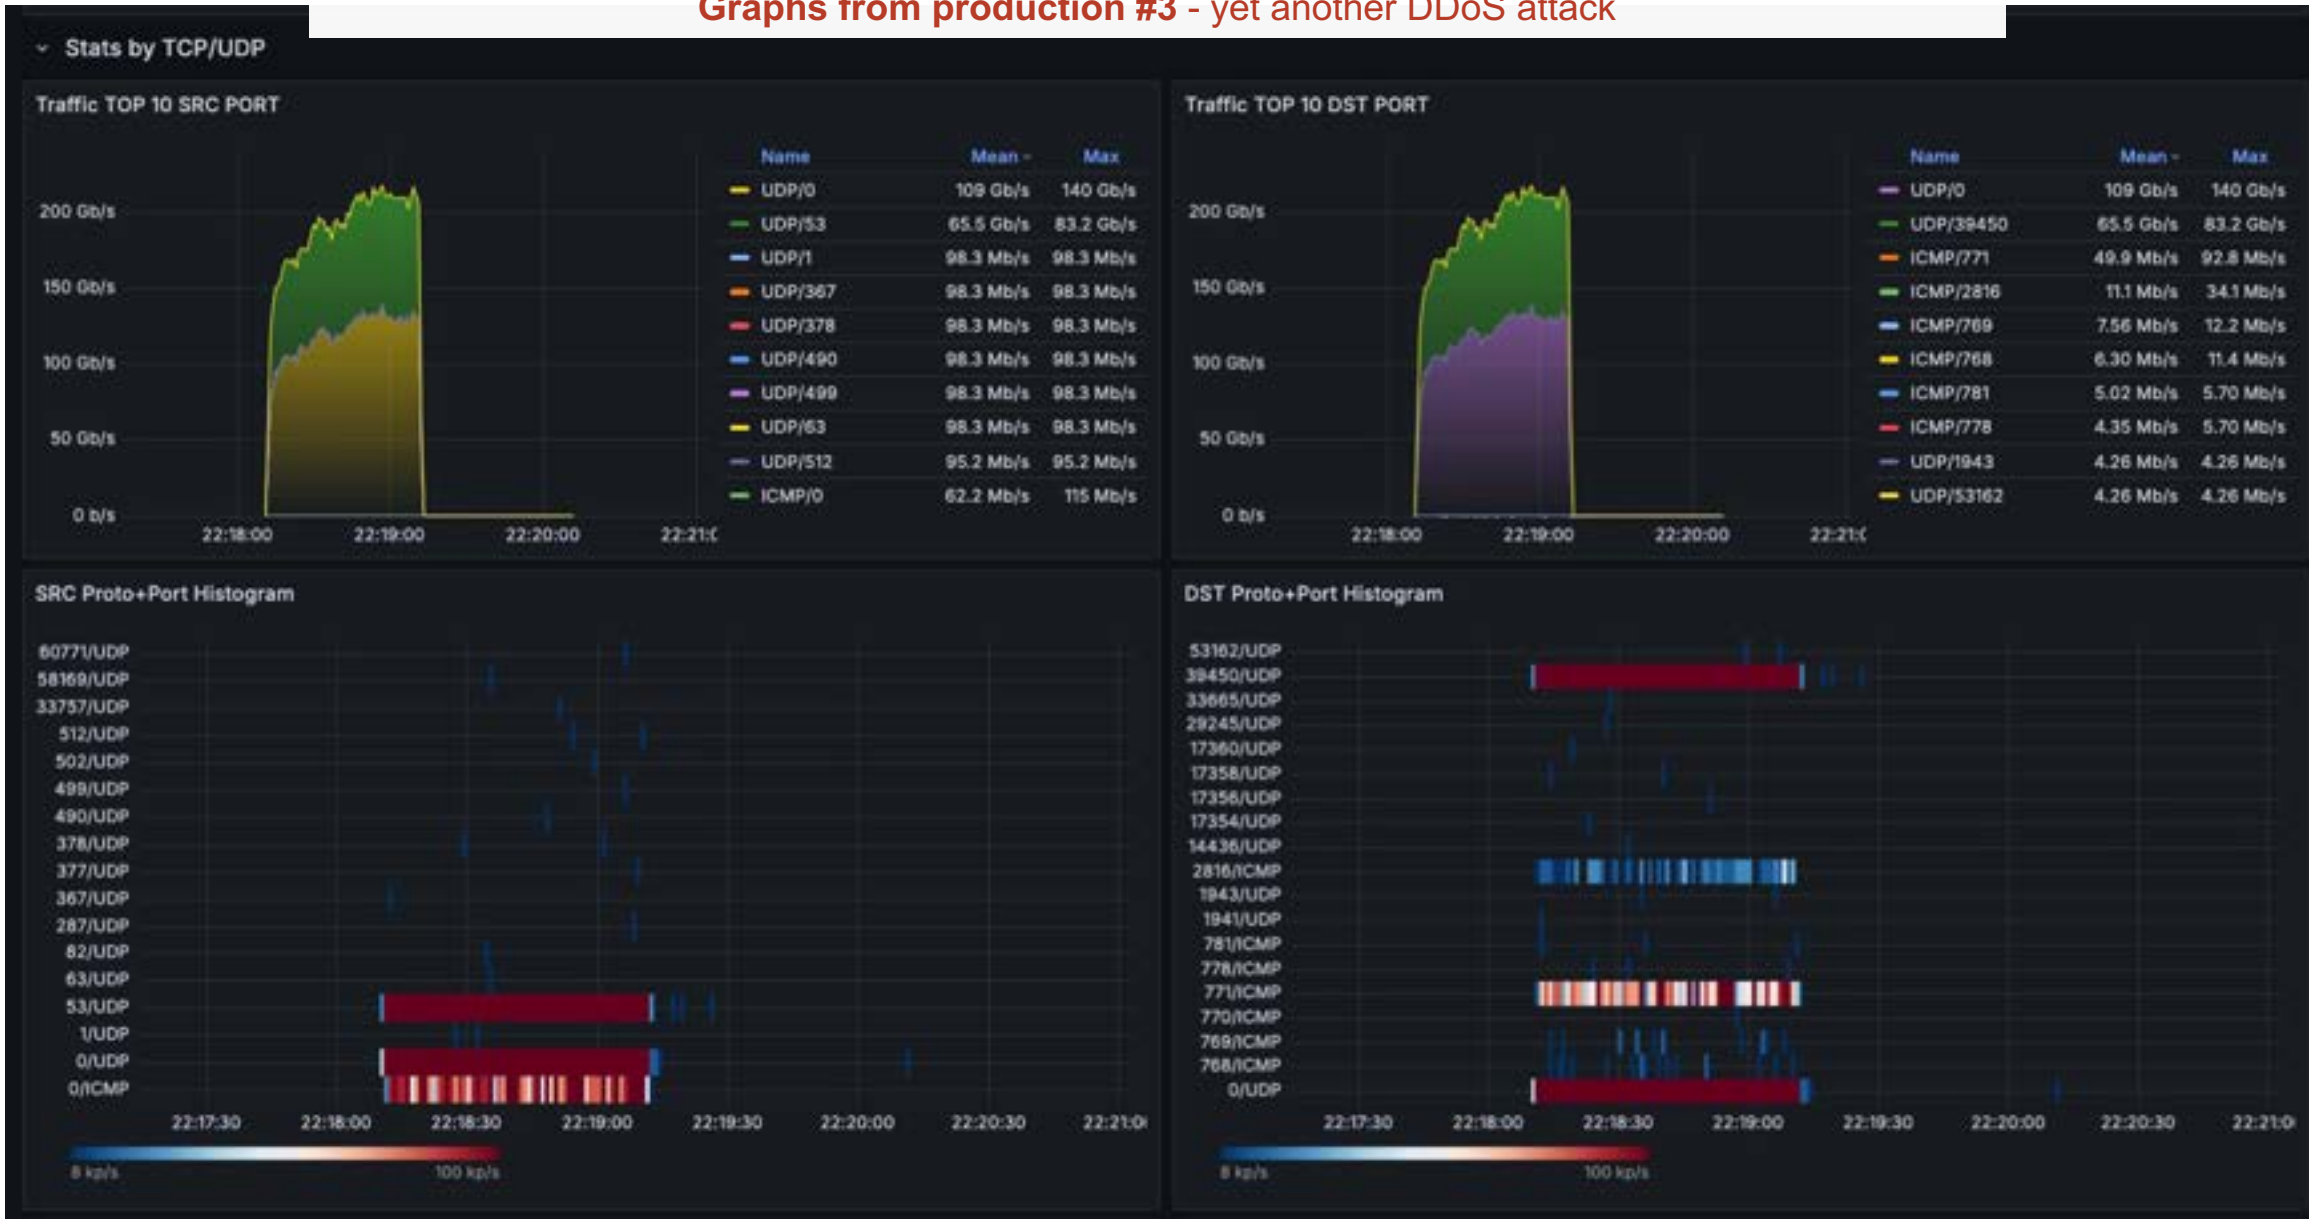
Graphs from production #2 - sleepless night, DDoS attack to customer



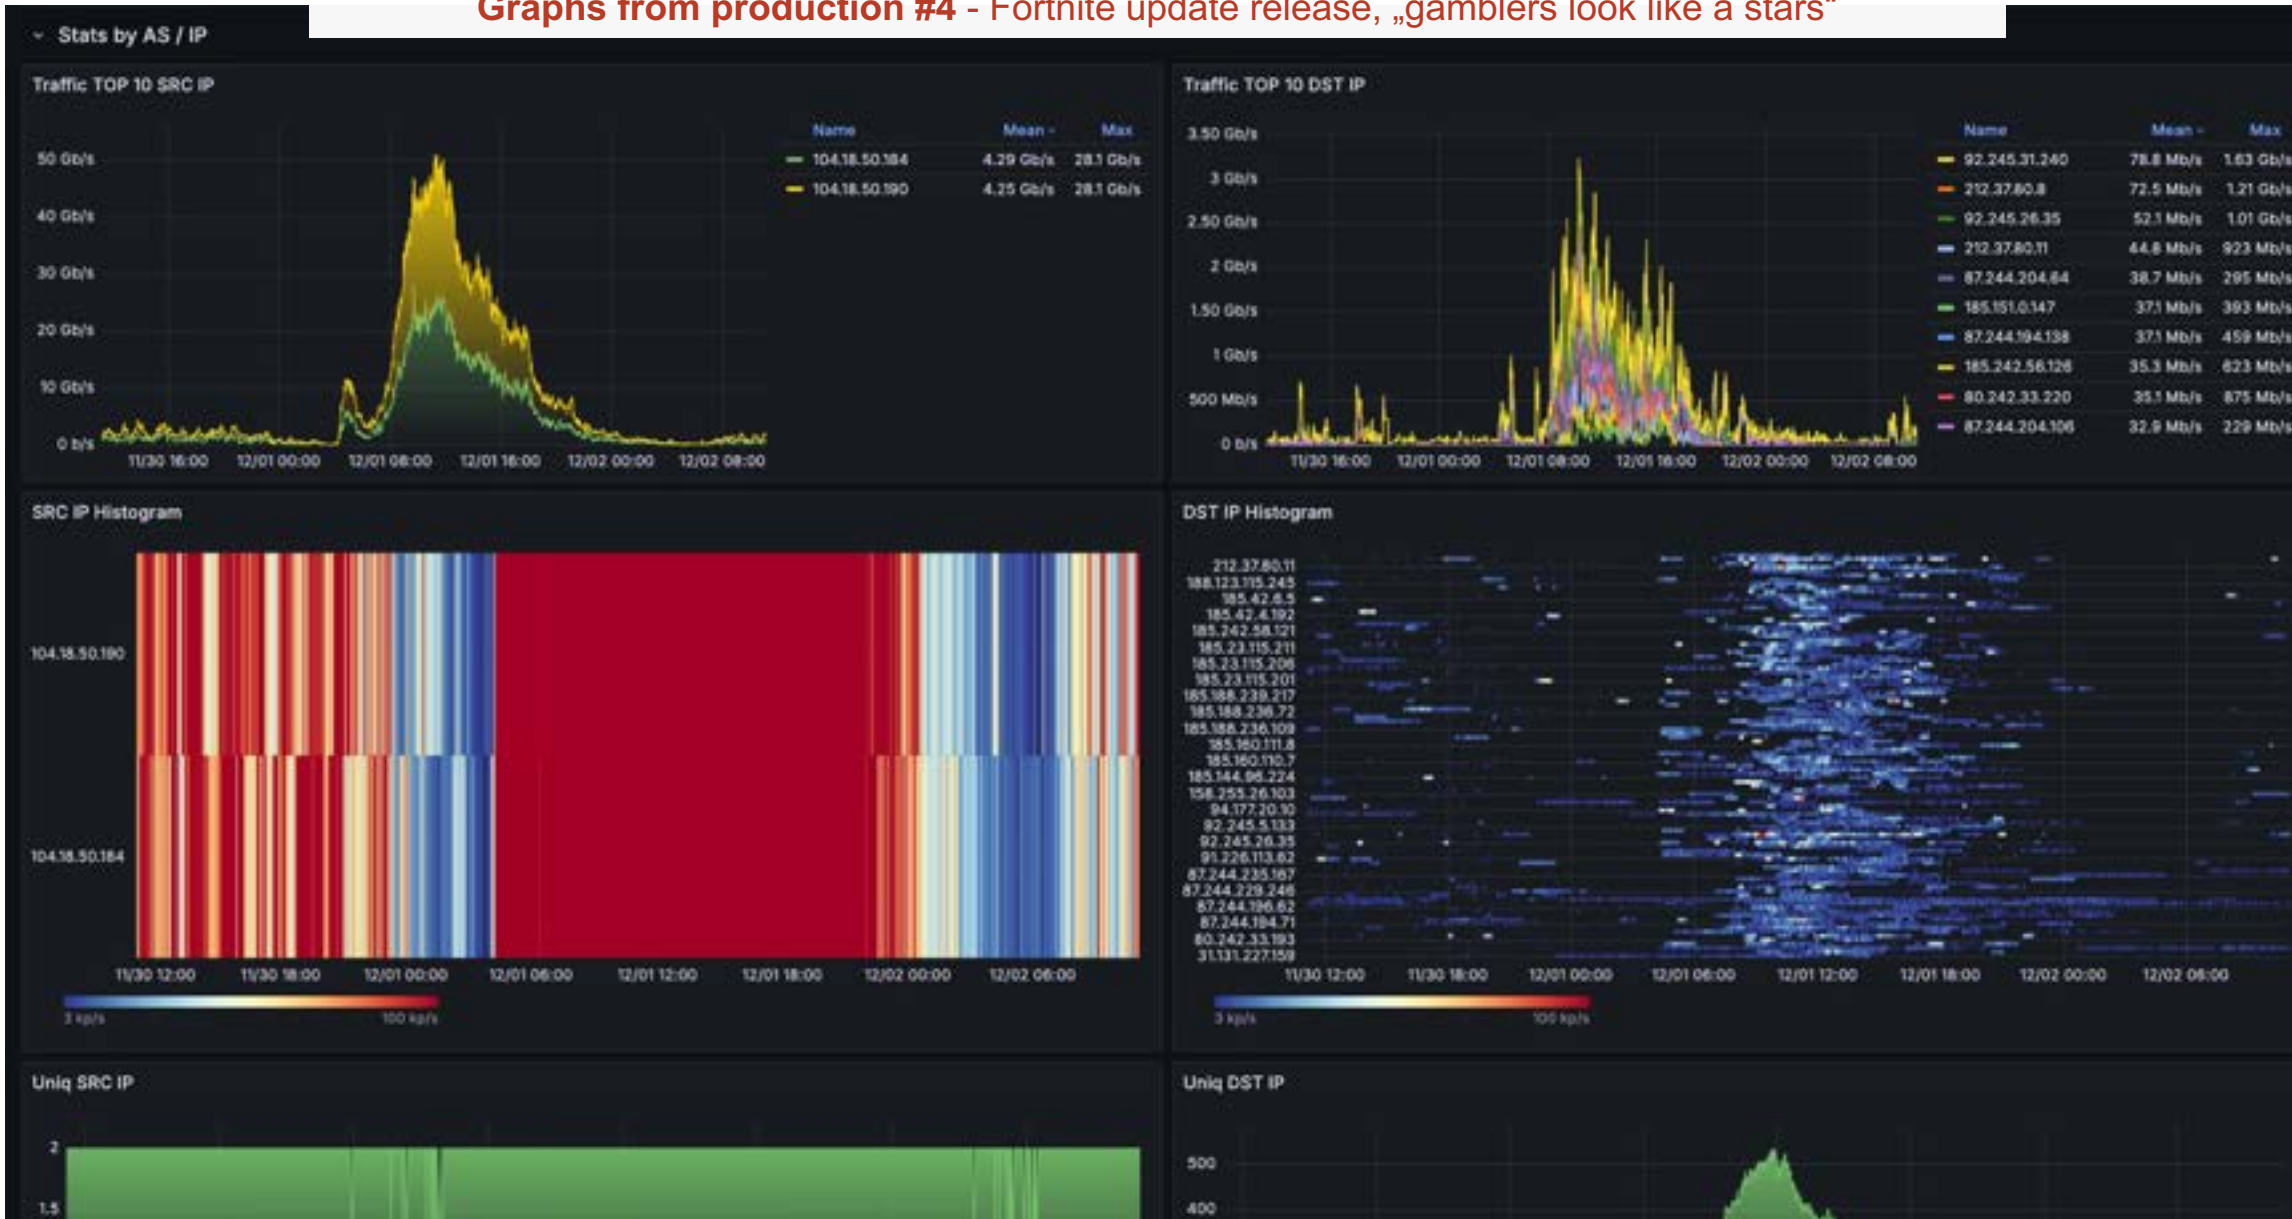
Graphs from production #3 - yet another DDoS attack



Graphs from production #3 - yet another DDoS attack



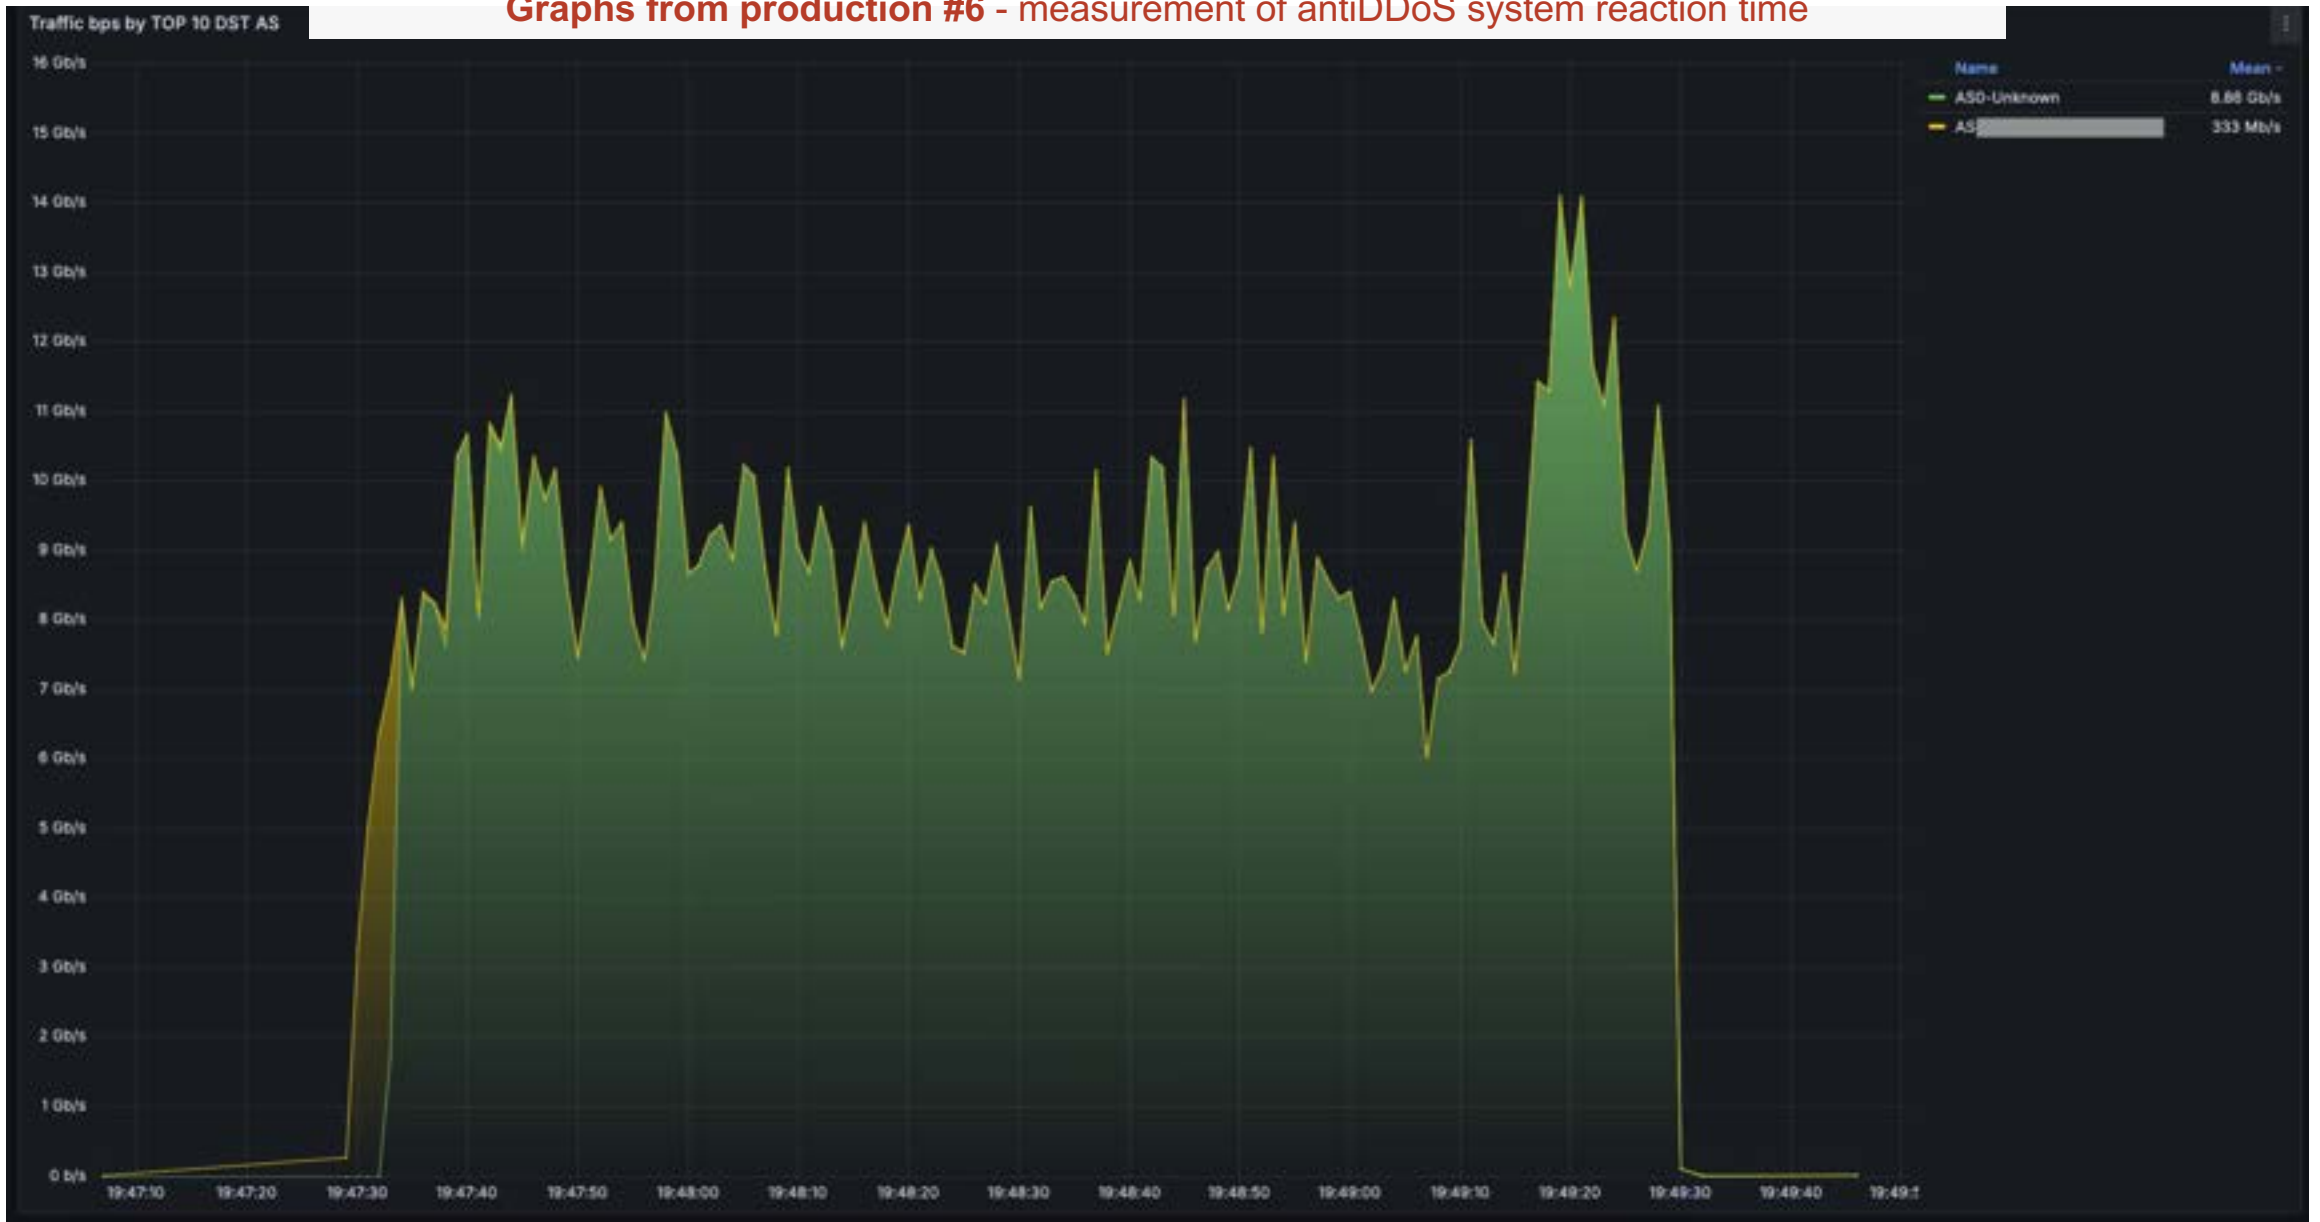
Graphs from production #4 - Fortnite update release, „gamblers look like a stars“



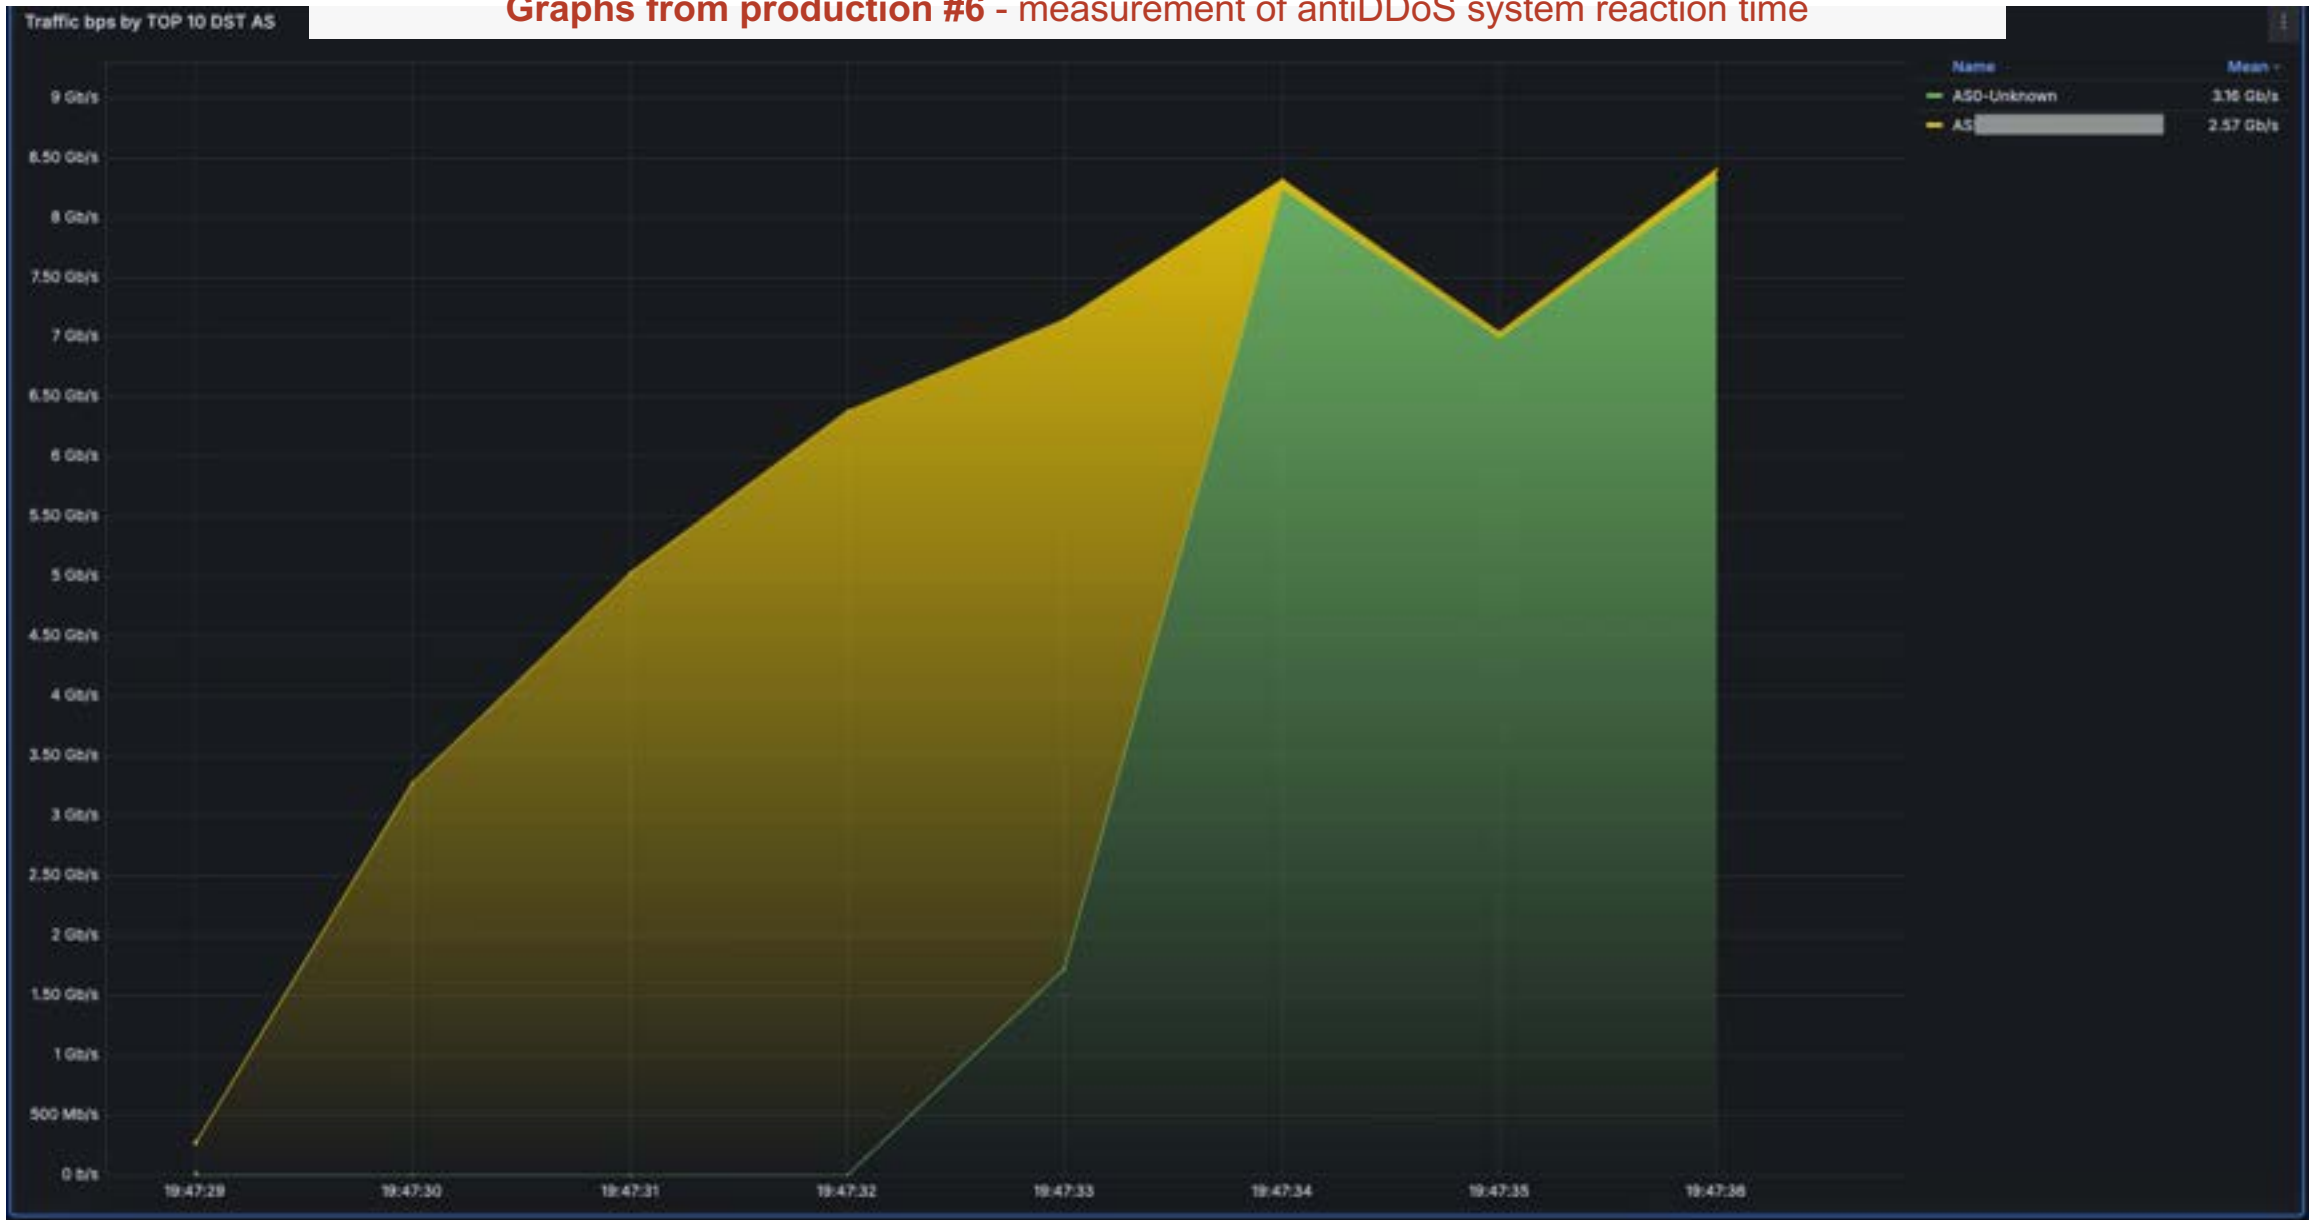
Graphs from production #5 - DDoS attack with destination port sliding each 7 seconds



Graphs from production #6 - measurement of antiDDoS system reaction time



Graphs from production #6 - measurement of antiDDoS system reaction time



A network diagram consisting of several red circular nodes connected by thin grey lines, forming a complex web of connections. The nodes are distributed across the left and top-left portions of the slide.

Thank you for your attention |

Blažej Krajňák
CSNOG | 22.1.2025