



Executive Summary

Kentik turns network traffic into operational and business value.

FOUNDED

2014

HQ

San Francisco

CUSTOMERS

200+

TEAM MEMBERS

70+

RUN BY

Network and
measurement nerds

GROWTH

20x since
January 2016

FOCUS

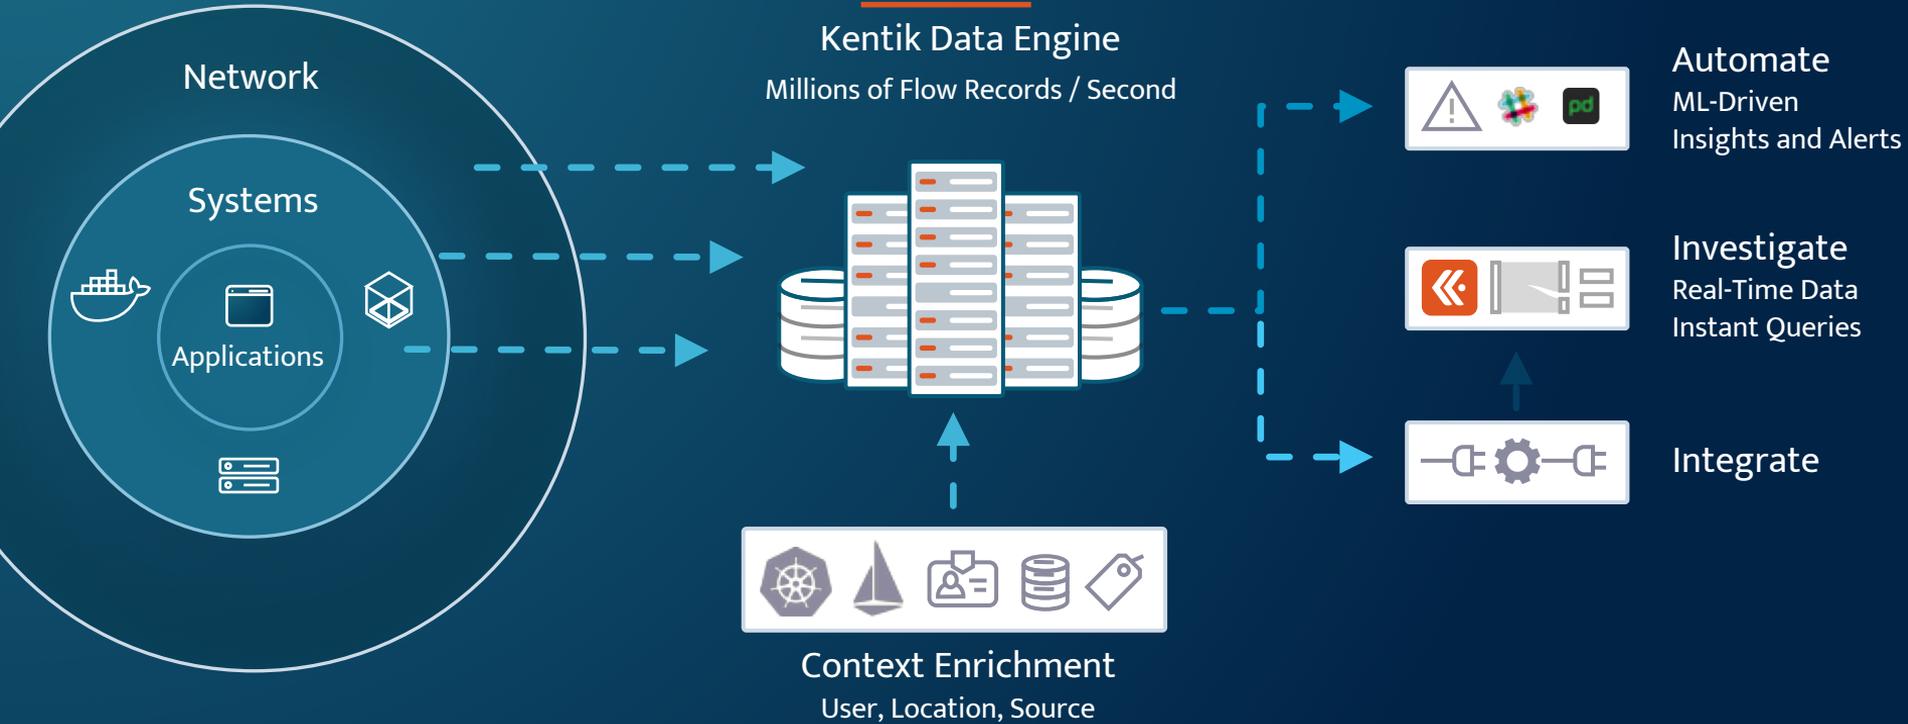
SPs and
Enterprise

TECHNOLOGY

In-house bg data platform
Delivered as a service



Kentik Platform



Kentik Platform

SOURCES

BUSINESS DATA

- Identity
- Orchestration
- OSS/BSS
- CRM, ERP
- Threat DBs

APP DATA

- kprobe
- Servers
- Containers
- Hypervisors
- CDNs
- DNS

NETWORK DATA

- Switches
- Routers
- Firewalls
- IoT

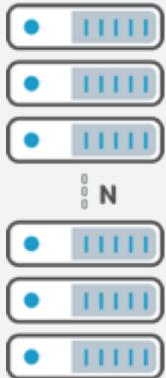
DATA

- Kentik Metadata API

- Host NPM
- Sensor NPM

- NetFlow
- VPC Flow Logs
- BGP
- SNMP

INGEST & FUSION

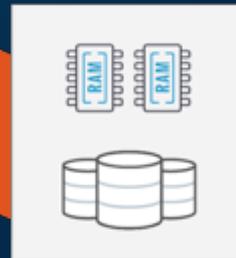


- GeoIP
- Kentik Global Threat Intel

STORAGE & QUERY



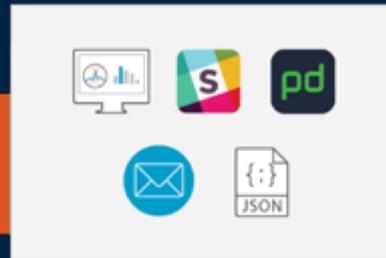
STREAMING & AGGREGATION



CLIENTS



ACTION TRIGGERS



Network Observability and Intelligence Use Cases

App-Aware Infrastructure Operations



Traffic
Engineering



Performance
Management

Application Operations



App vs.
Network
Performance



Application
Governance

Attack Detection, Mitigation & Investigation



Threat
Detection



DDoS Defense



Digital
Forensics

Business Operations



Service
Creation



Cost
Management



Adding context by
enriching traffic data



What do we want?

Enrichment:

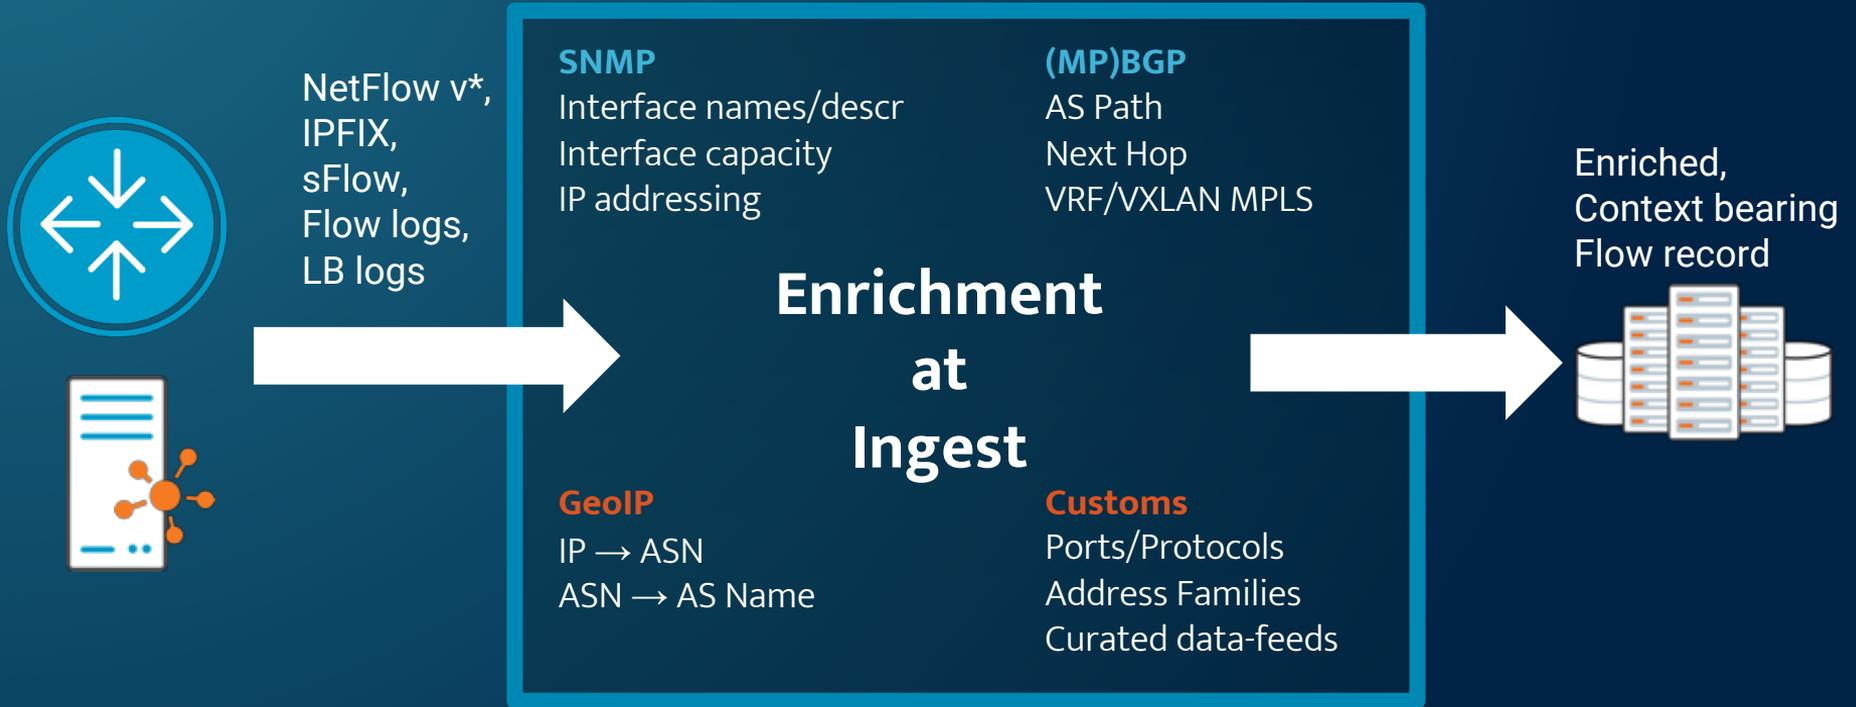
Going beyond the basics, and
contextualizing network data.

How do we do it?

Tags and tables:

Dynamic routing tables of context -
what does the traffic *mean*?

Making Data Useful: Basic Flow Enrichment



Next Up: Infrastructure Context

1st Class Citizen / Built-In

Interface Classification

Network Classification

Custom AS Groups

Customer/Provider tagging

Custom Geo

Overlay Service, Threat Data

Clouds

Threat feeds

Applications tagging

Application

k8s, Orchestration

CDN Logs

Istio, nginx, Load
Balancer

Next Up: Infrastructure Context

1st Class Citizen / Built-In

Interface Classification

Inside/outside directionality

Connectivity type

Provider vs Customer

Network Classification

e2e directionality

Customer/Provider tagging

CRM meets flows

Custom AS Groups

Networks w/ multiple ASNs

Private ASNs

Custom Geo

Country groups/Markets

Sub country groups

Overlay Service, Threat Data

Clouds

ISP Embedded + Self-hosted CDNs

Cloud providers

Threat feeds

Botnets

Infected hosts

Applications tagging

OTT services

Well known Apps

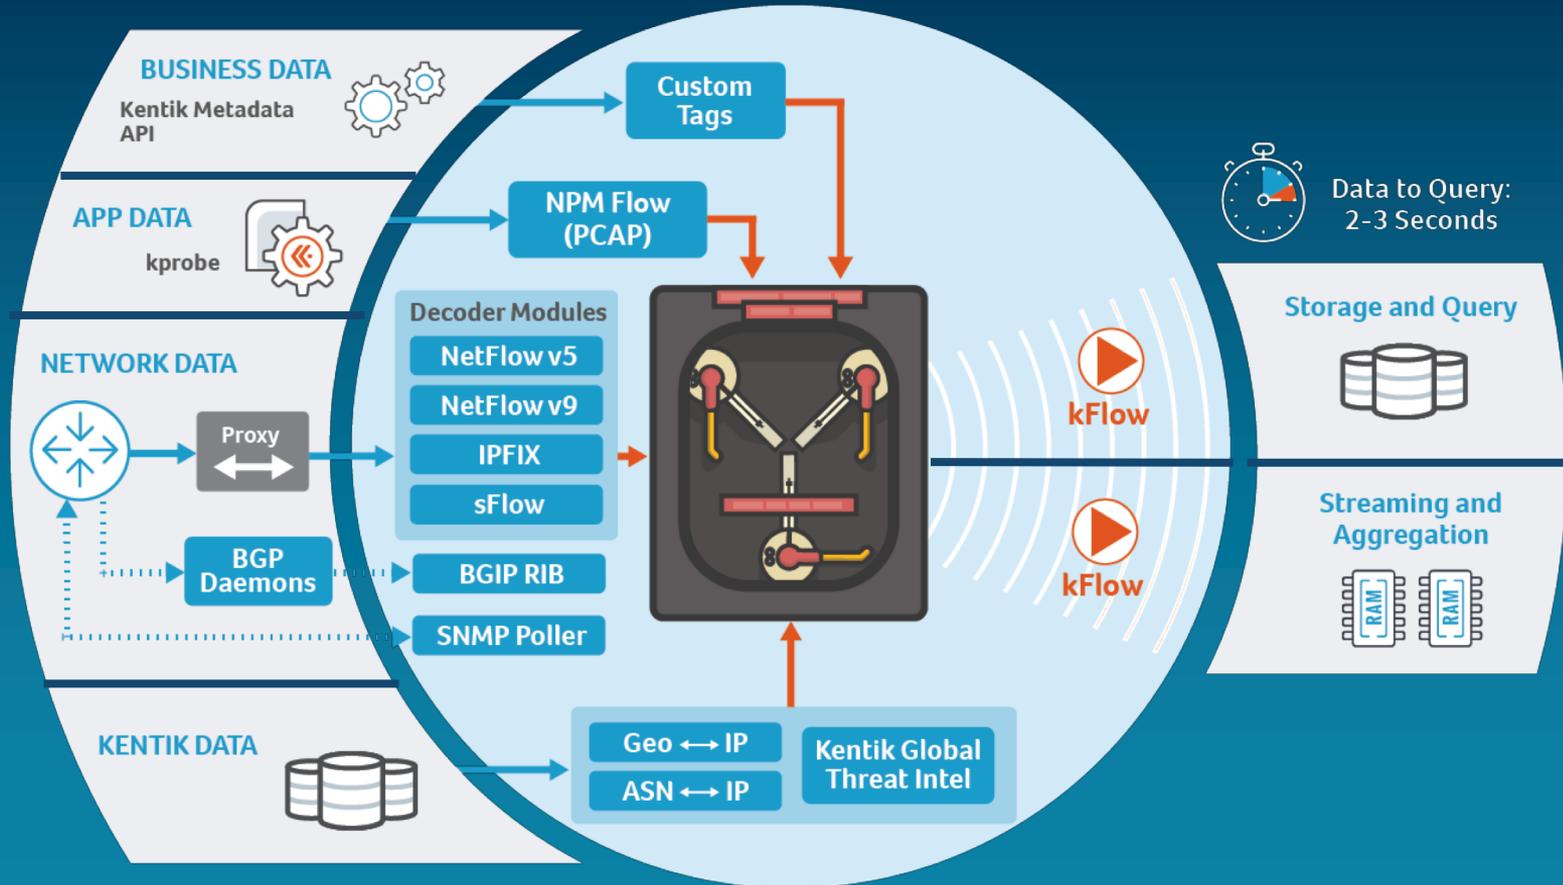
Application

k8s, Orchestration

CDN Logs

Istio, nginx, Load
Balancer

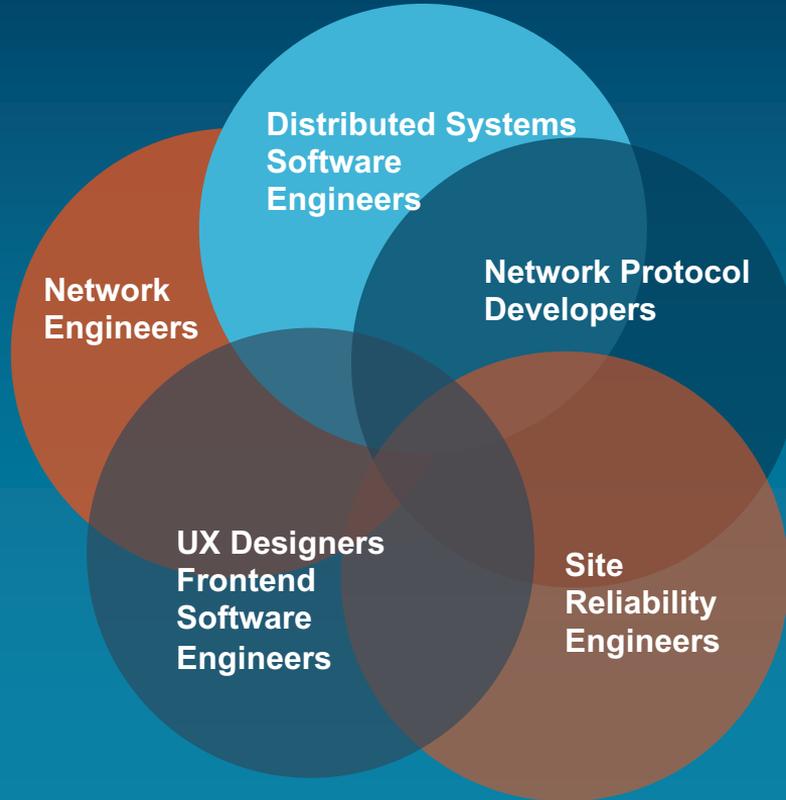
Data Fusion



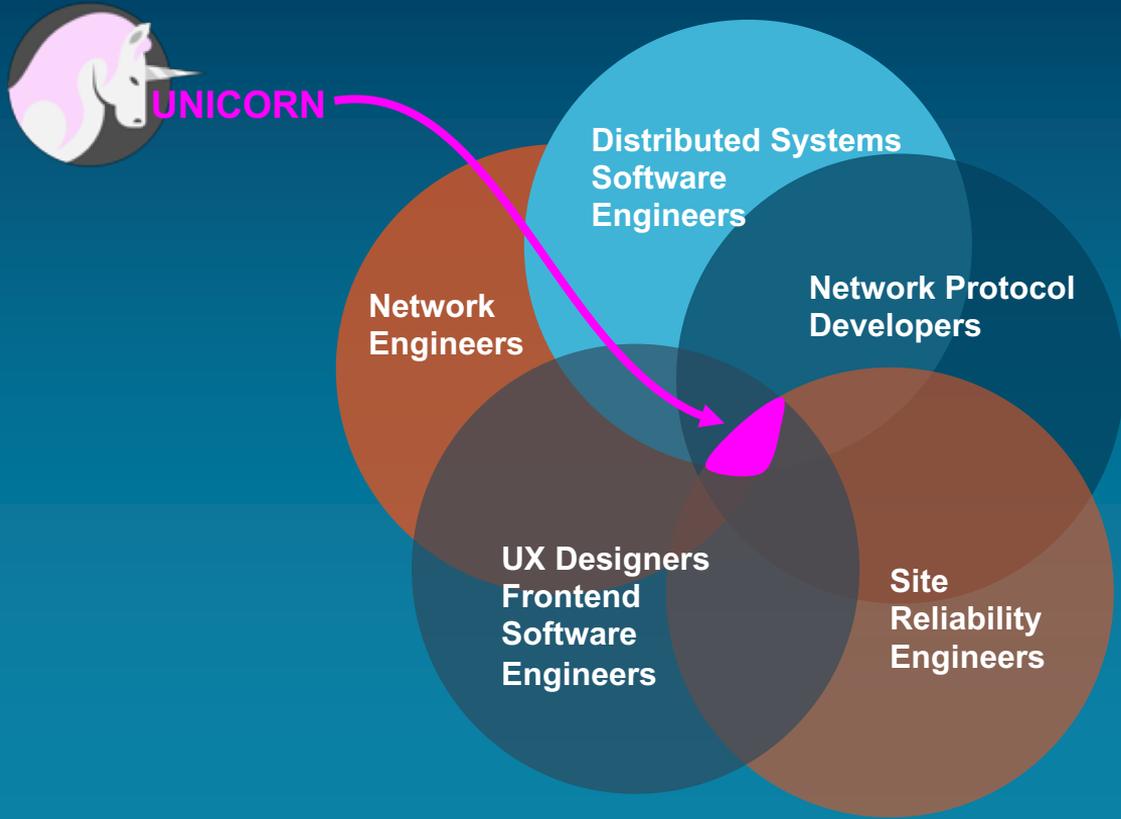
The How

- LOTS of routing tables!
- Can be IP, MAC, VLAN, device, interface, BGP, or other traffic field-based
- Up to dozens of tables with millions of entries, per customer
- Live updating in real time through ingest system
- With persistence
- Tables must be synced with load balancing
- Includes global tables of Geo, threats
- In production - millions of tags * millions of FPS, ~20ms avg update

Enrichment requires cross-disciplinary skills



How to make more unicorns?



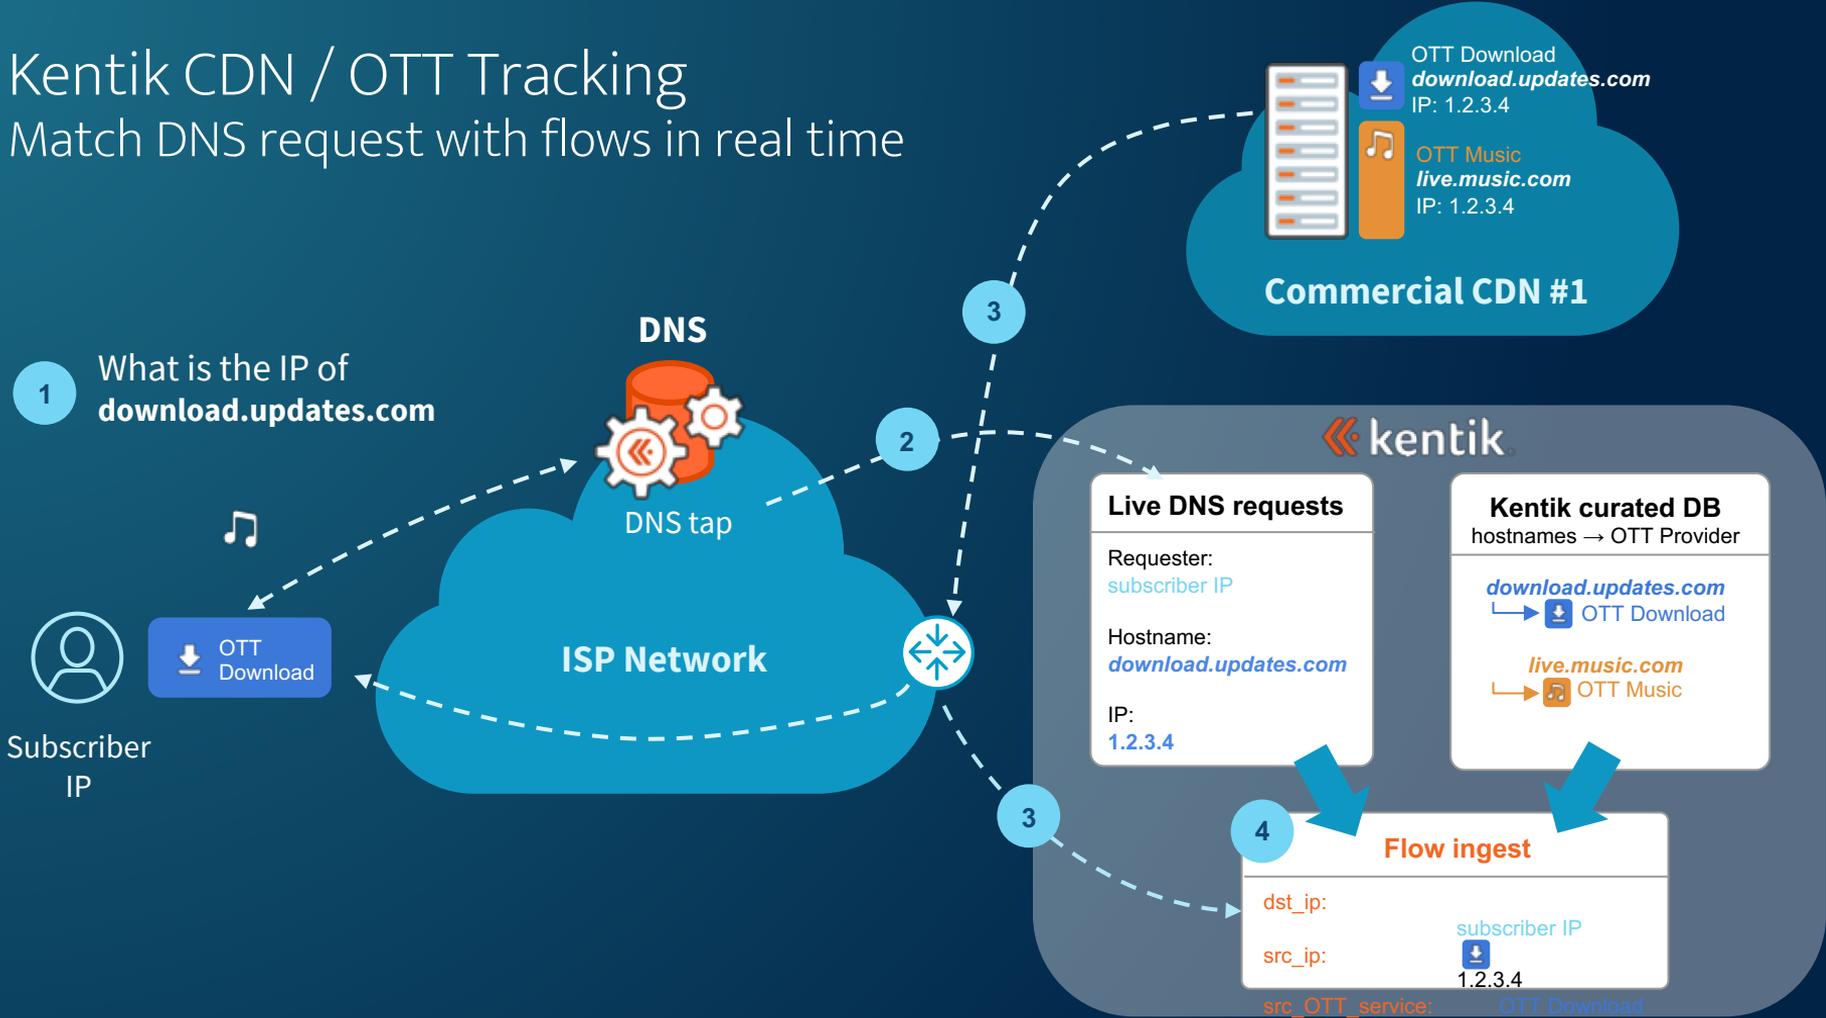


What does enrichment enable?



Kentik CDN / OTT Tracking

Match DNS request with flows in real time



BGP Ultimate Exit

Examining traffic exit points on your network

Peer or
Customer

BGP

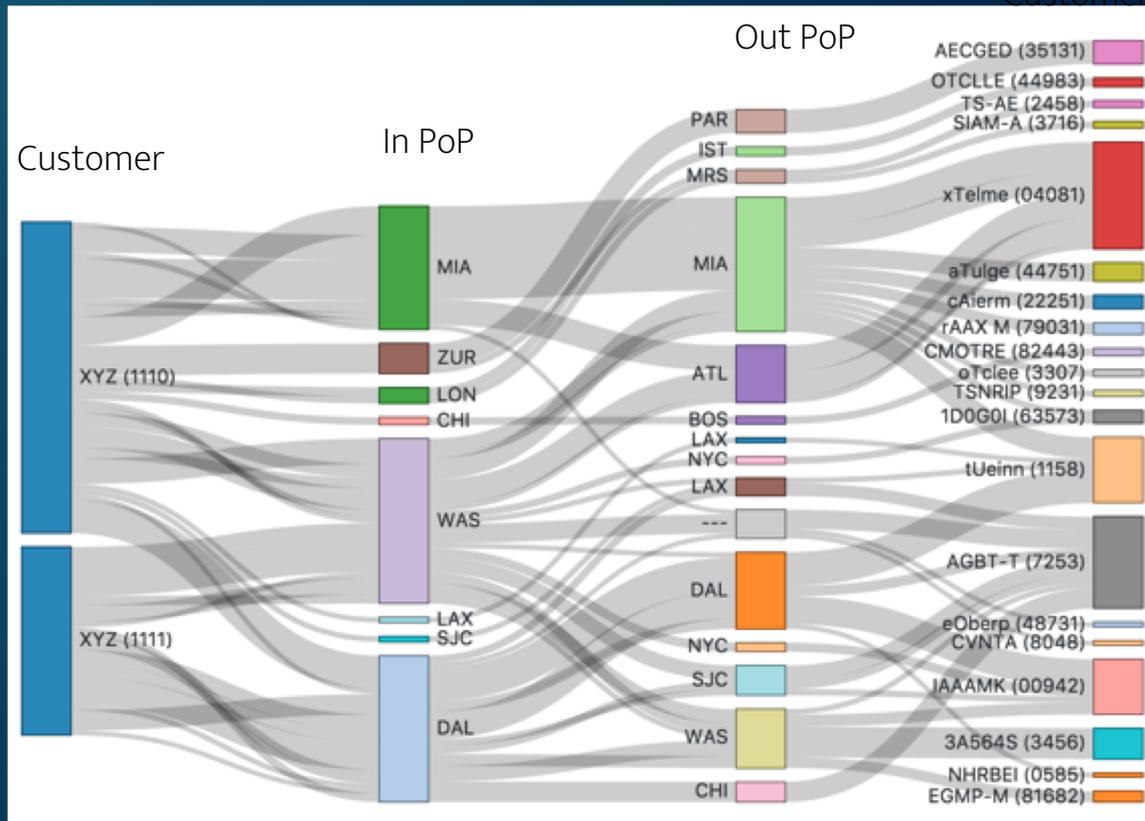
All flow tools allow you to view where traffic *enters* your network, can your tool show you where it *leaves*?

BGP-UE uses iBGP data from your devices to determine where flows exit, allowing you to get a deeper understanding of your traffic.

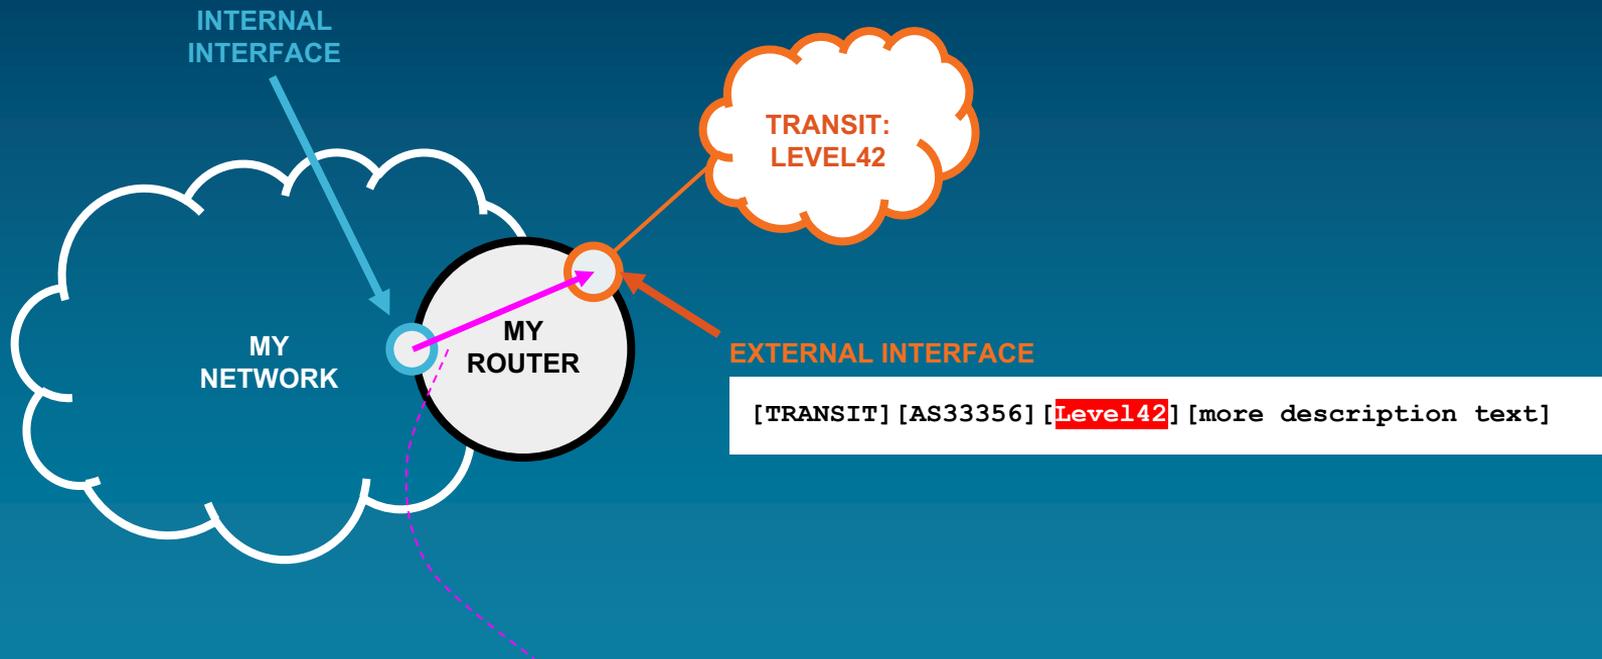
Is my CDN customer keeping traffic local?

Is my transit customer using expensive subsea capacity?

Find out with **Ultimate Exit**.



Enrichment on Traditional Networks: Interfaces on Devices



Enriched flow record

```
src_int: {INTERNAL, BACKBONE, n/a}
dst_int: {EXTERNAL, TRANSIT, LEVEL42}
```

Useful Enrichment: Interface Classification

INTERFACE DESCRIPTION (SNMP)

```
[TRANSIT] [AS33356] [Level42] [more description text]
```

DESCRIPTION MATCH REGEX (Enrichment engine)

```
^\[TRANSIT\]\[.*\]\[(.*)\].*$
```

INTERFACE CLASSIFIERS

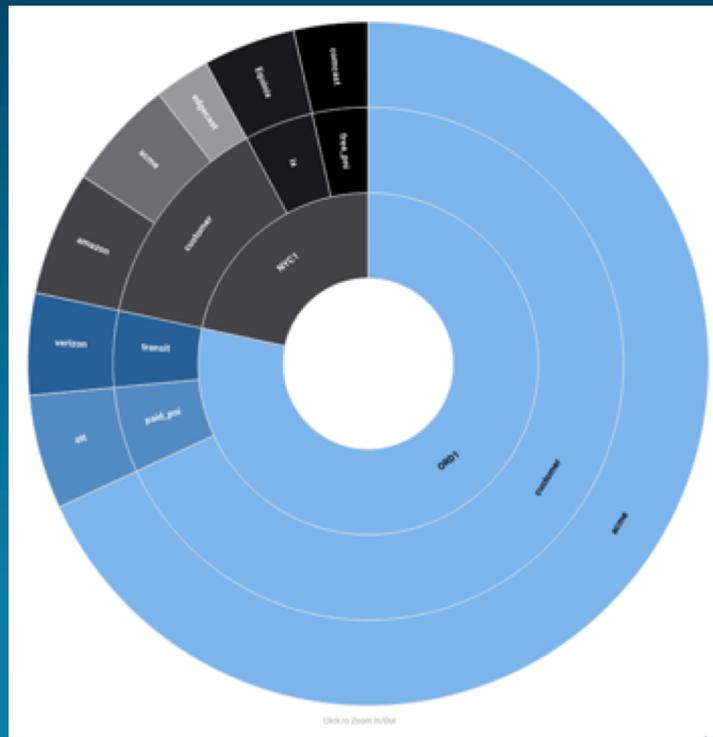
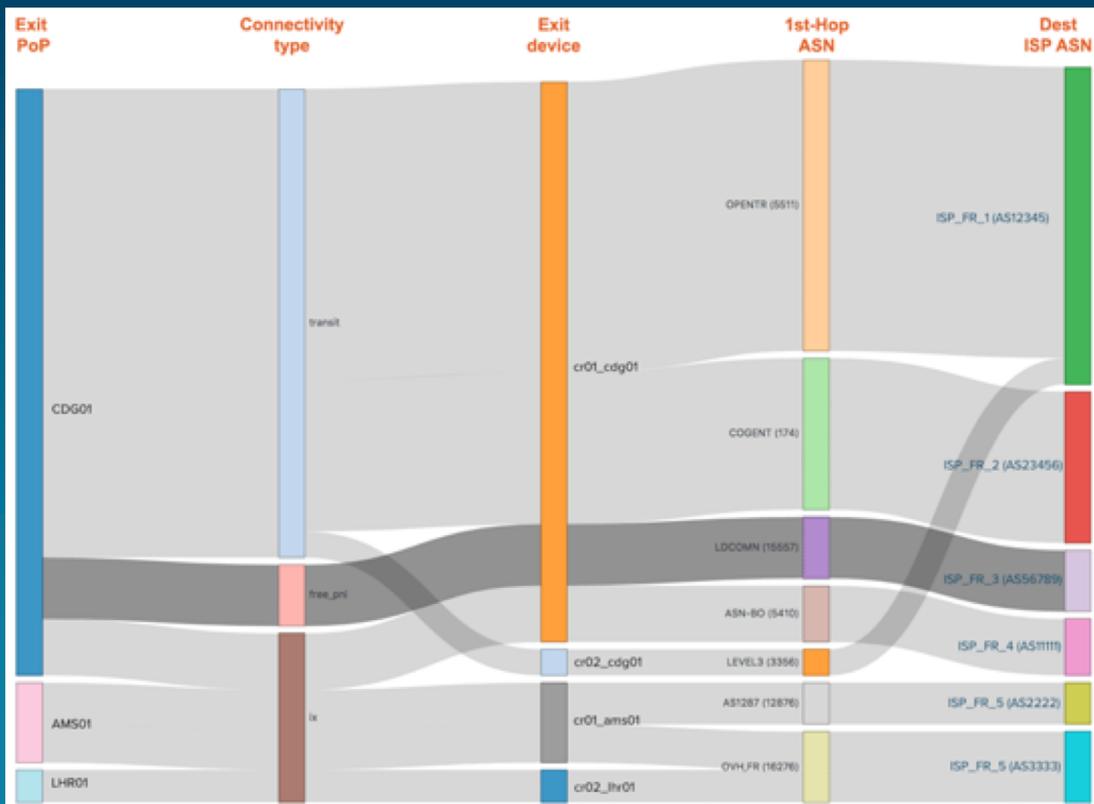
```
SET INTERFACE NETWORK BOUNDARY: EXTERNAL  
SET INTERFACE CONNECTIVITY TYPE: TRANSIT  
SET INTERFACE PROVIDER:  
(LEVEL42)
```

```
$1
```

Enriched flow record

```
src_int: {INTERNAL, BACKBONE, n/a}  
dst_int: {EXTERNAL, TRANSIT, LEVEL42}
```

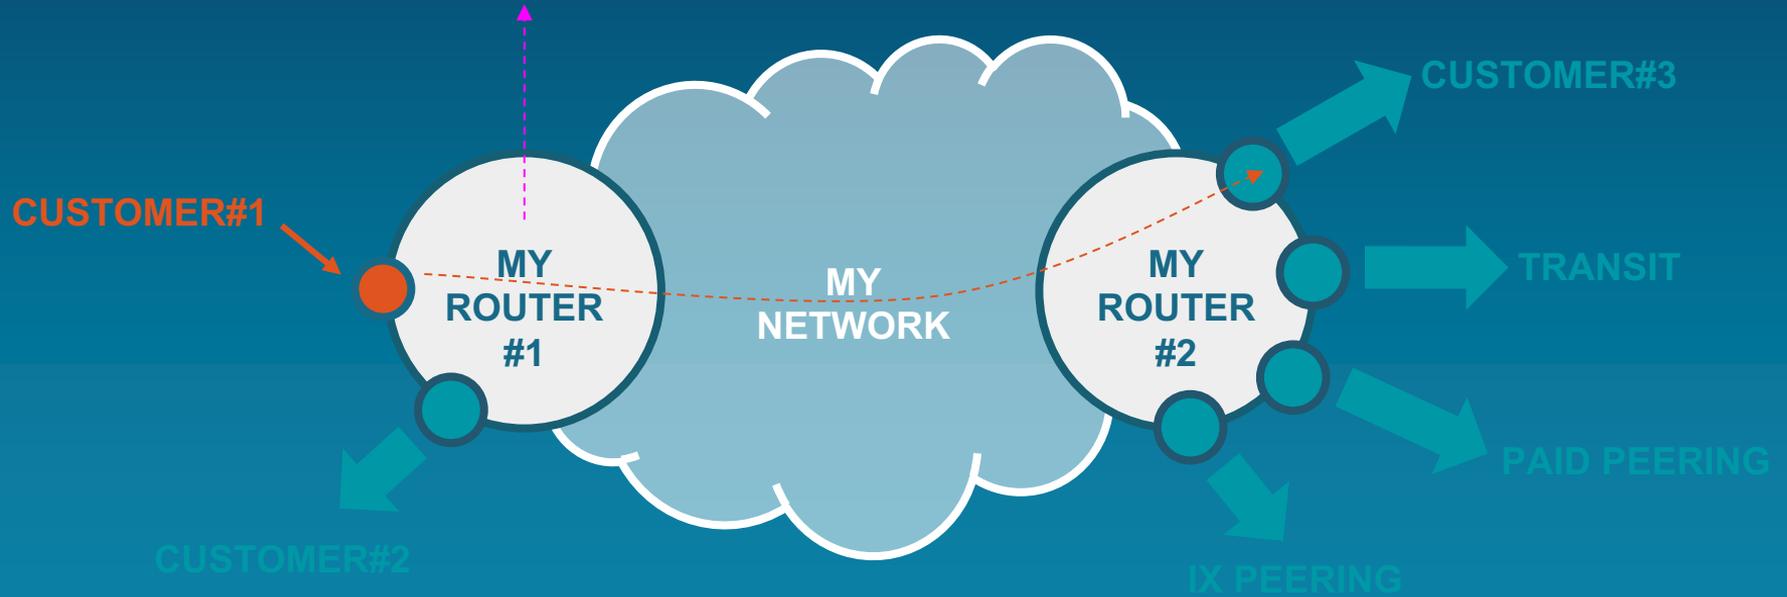
Useful Enrichment: Interface Classification



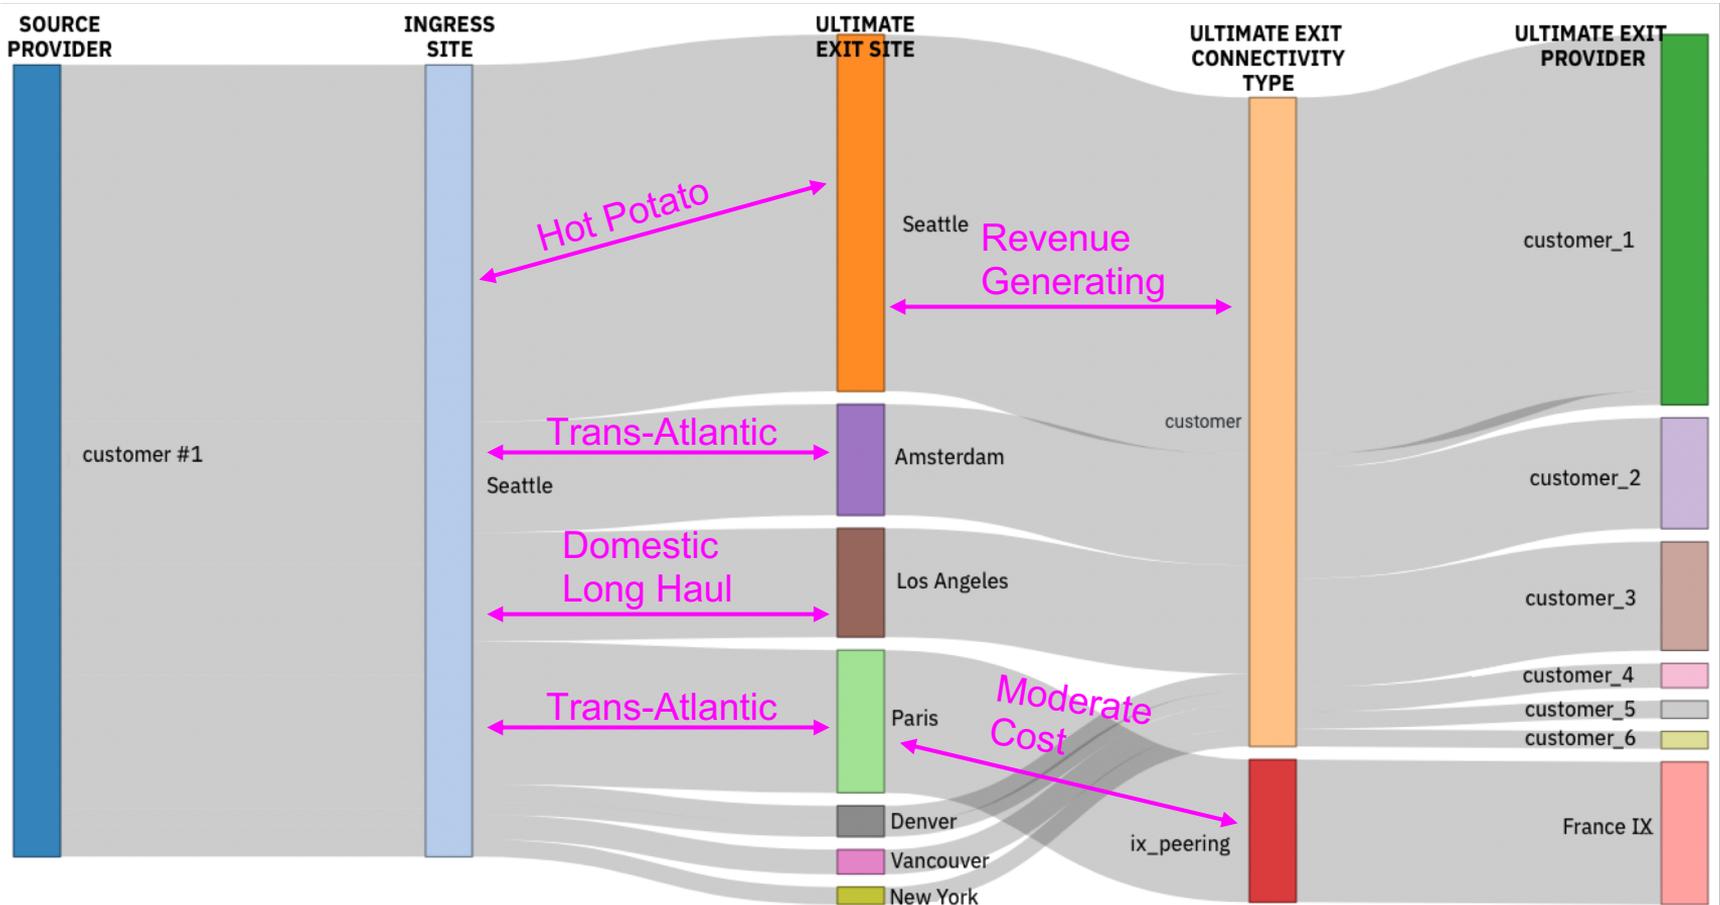
Ultimate Exit Discrimination

FLOW RECORD:

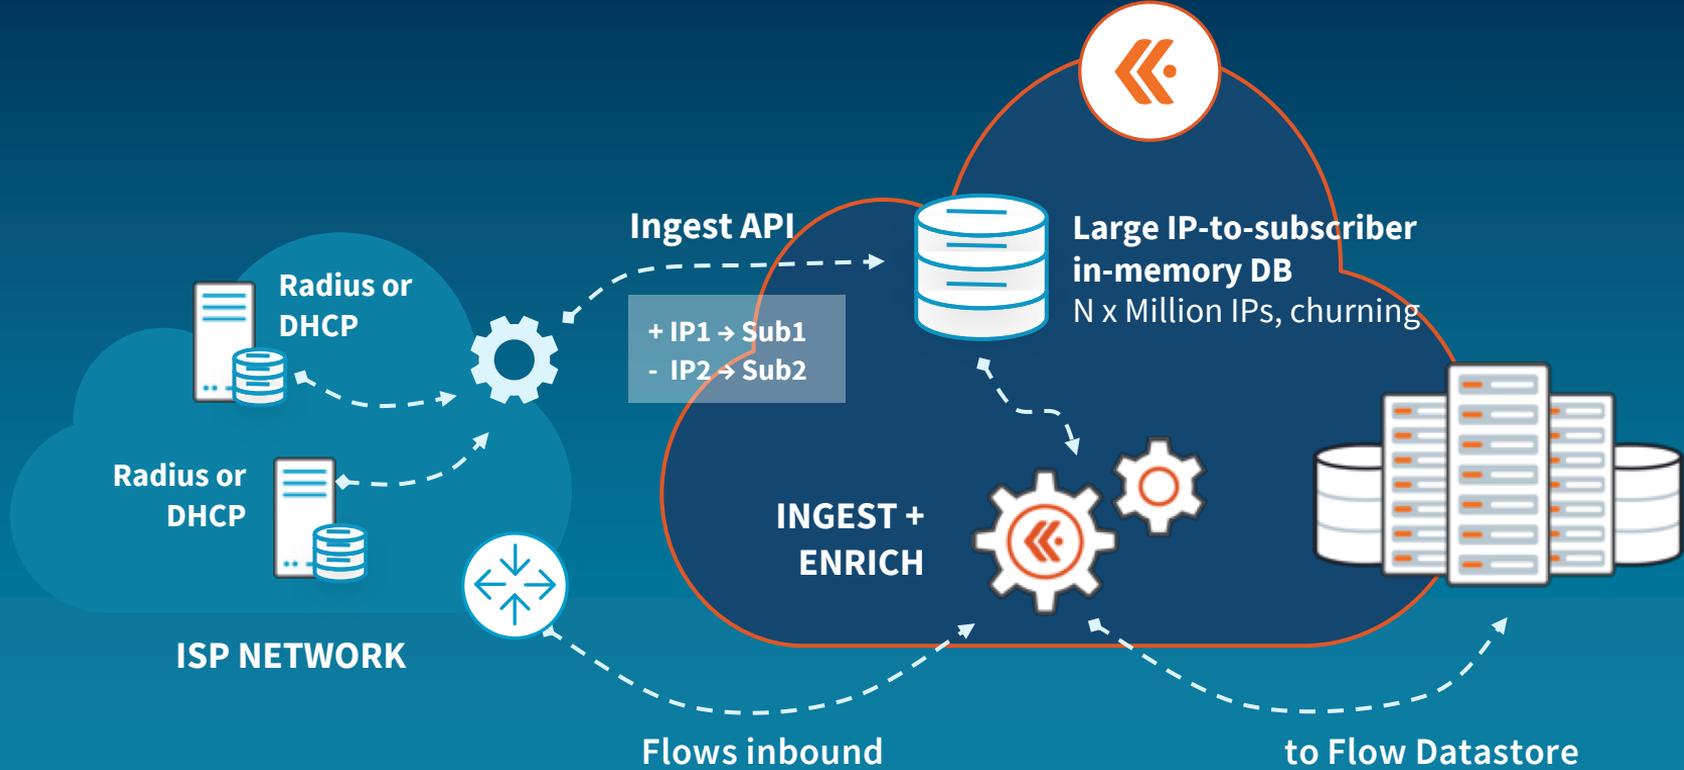
- Ultimate Exit {country, site, device, interface,}: {country,site, MyRouter#2, customer#3}
- Ultimate Exit Connectivity Type: customer
- Ultimate Exit Connectivity Provider/Customer: CUSTOMER#3



Ultimate Exit Discrimination

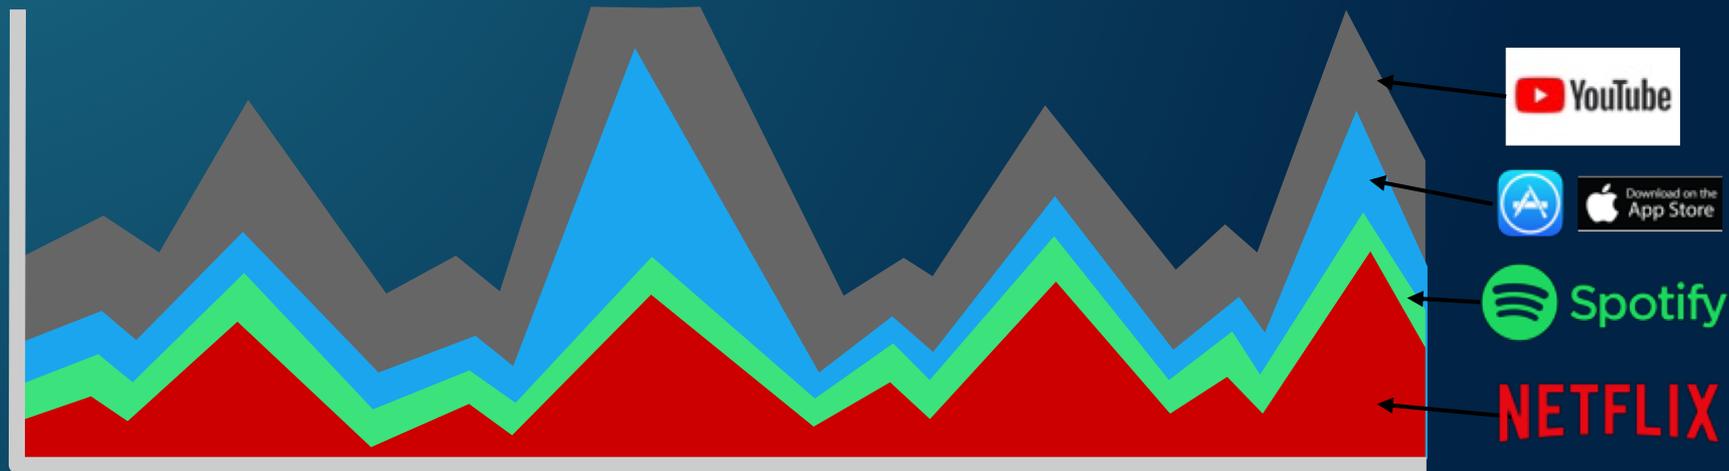


Subscriber Tracking: Implementation



OTT Service Tracking

Mark flows with Src/Dst OTT services being used



CDN Attribution

Discovering sources of traffic in neighboring networks

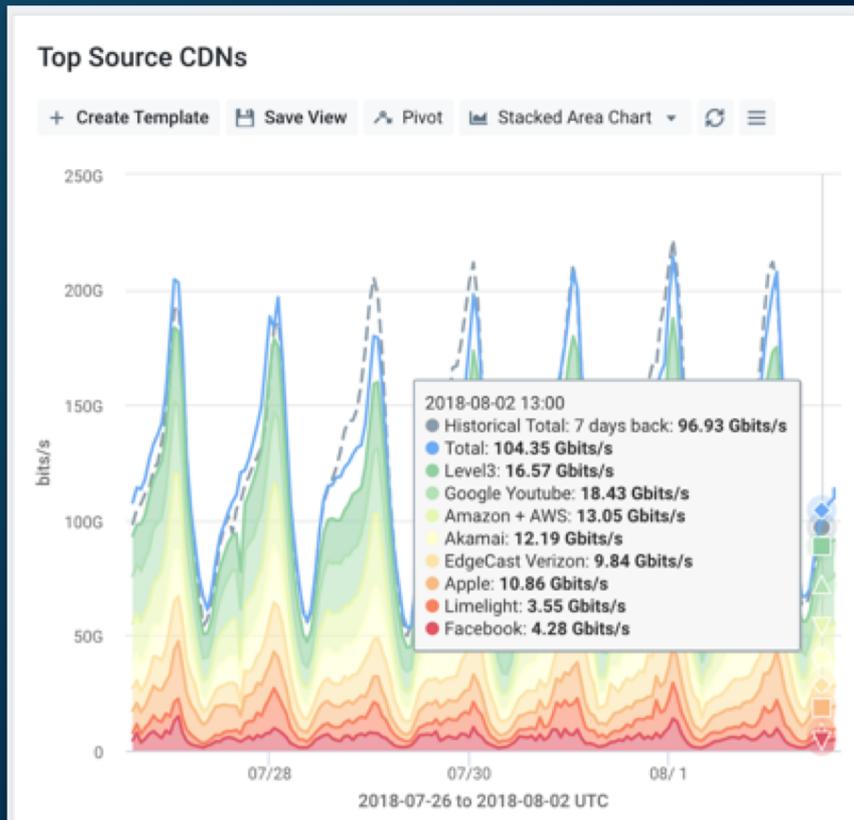
CDN

Is all the traffic from my peer actually from their *customers* or from embedded *CDNs*?

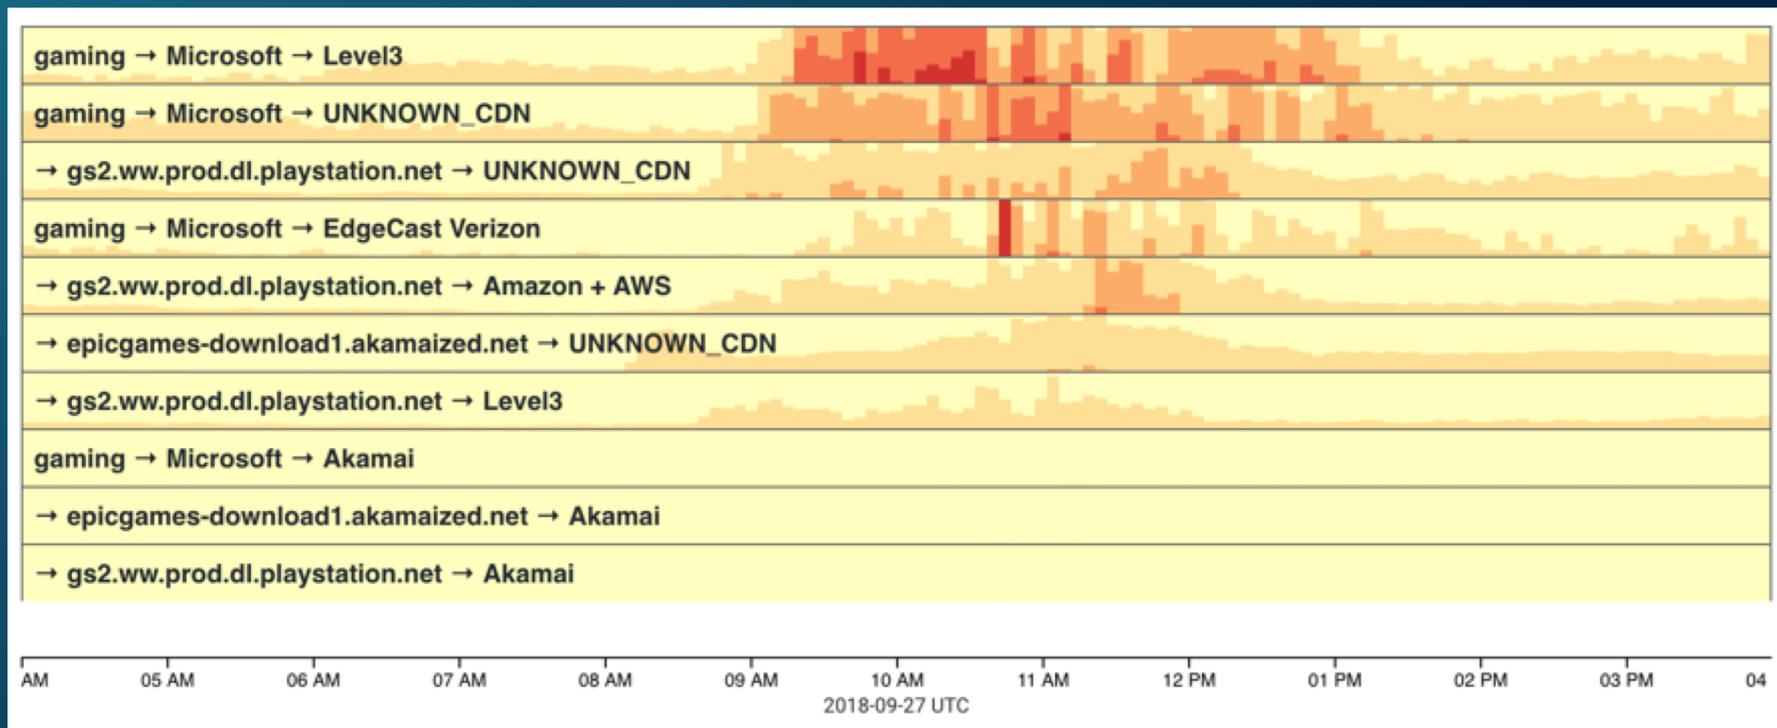
Using DNS data extracted from your network, we can determine if your neighboring ASNs have embedded CDN caches and what percentage of the overall traffic they represent.

Is my CDN customer delivering traffic locally?

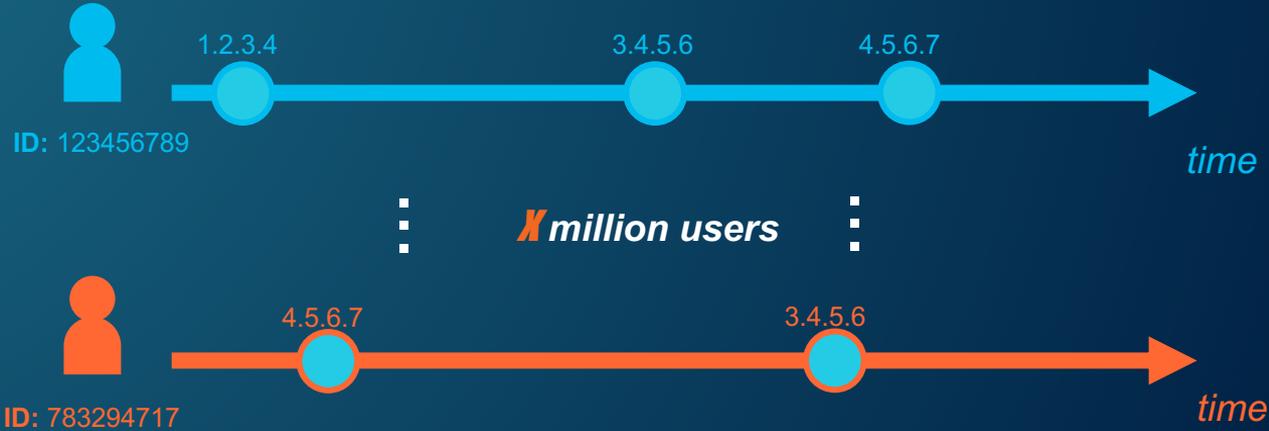
Find out with **CDN Attribution**.



OTT Tracking in Action



Use Case: Subscriber Tracking



Simple tools for high-tier customer support

“UserID 1234 peaks at XX Mbps, let’s compare to their plan”

Fair-Use quota monitoring

“User 4567 complains about excess usage charges, which apps / services are responsible?”

Validate engineering assumptions on user bandwidth consumption

“Users on this CMTS/DSLAM/PoP consume in 95%ile YY Mbps at peak time”



Thanks!

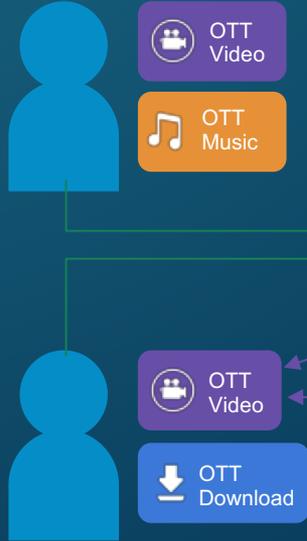
kentik.com/nfd19



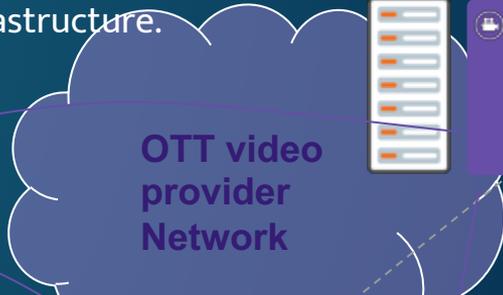
OTT Tracking Challenges

CDNs vs Owned vs Embedded delivery infrastructure.
CDNs host multiple OTT providers.

ISP Subscribers



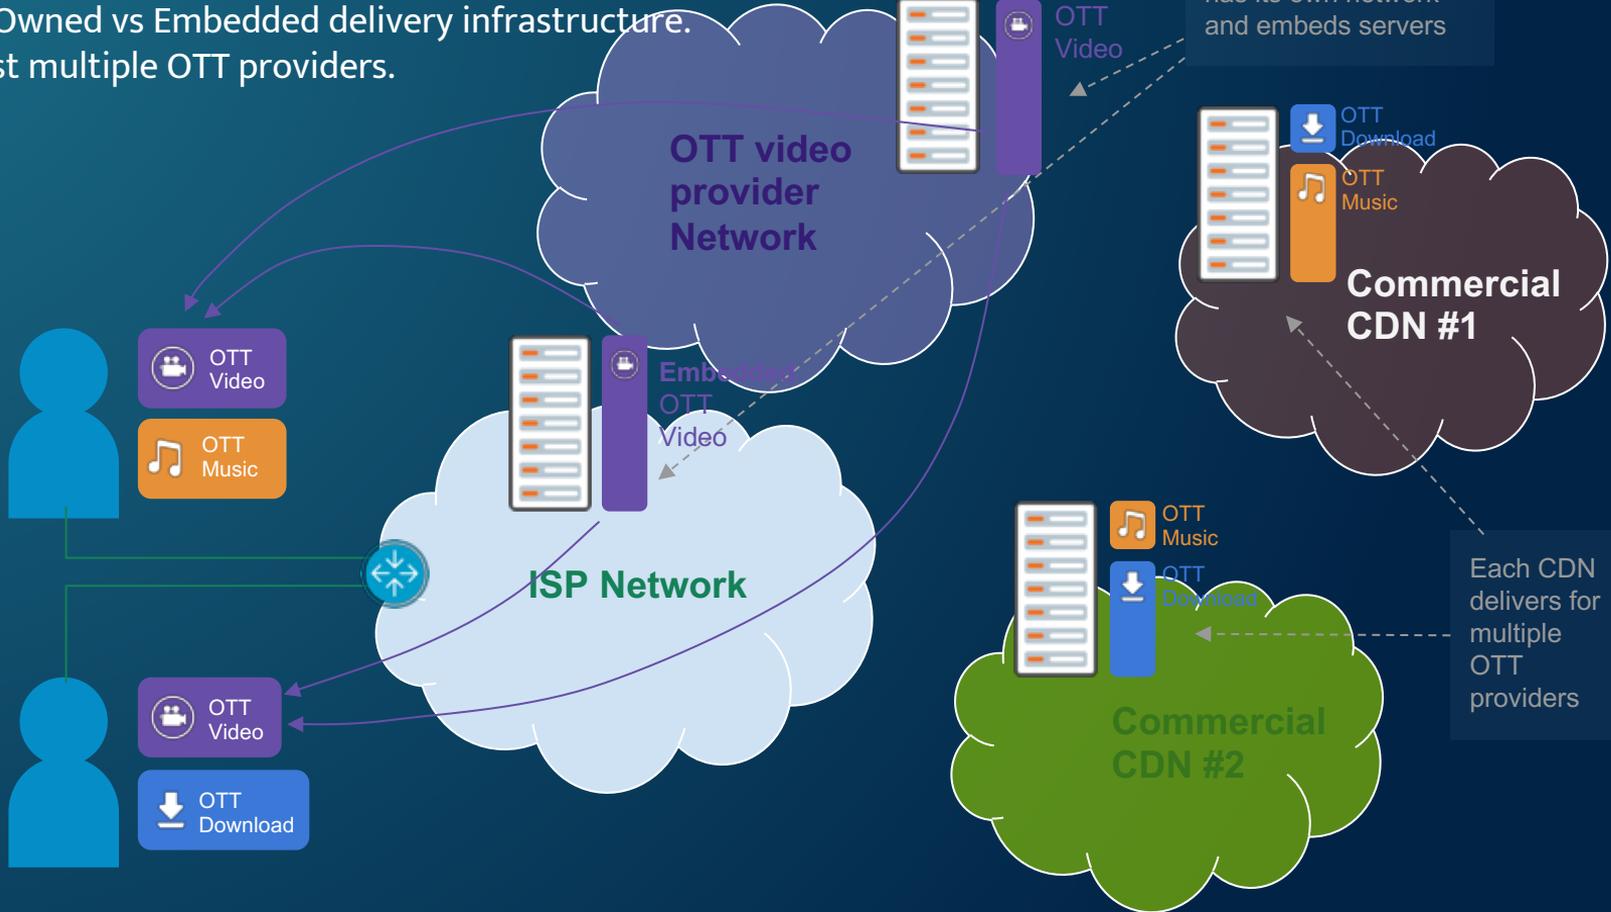
ISP Network



OTT Video Provider has its own network and embeds servers

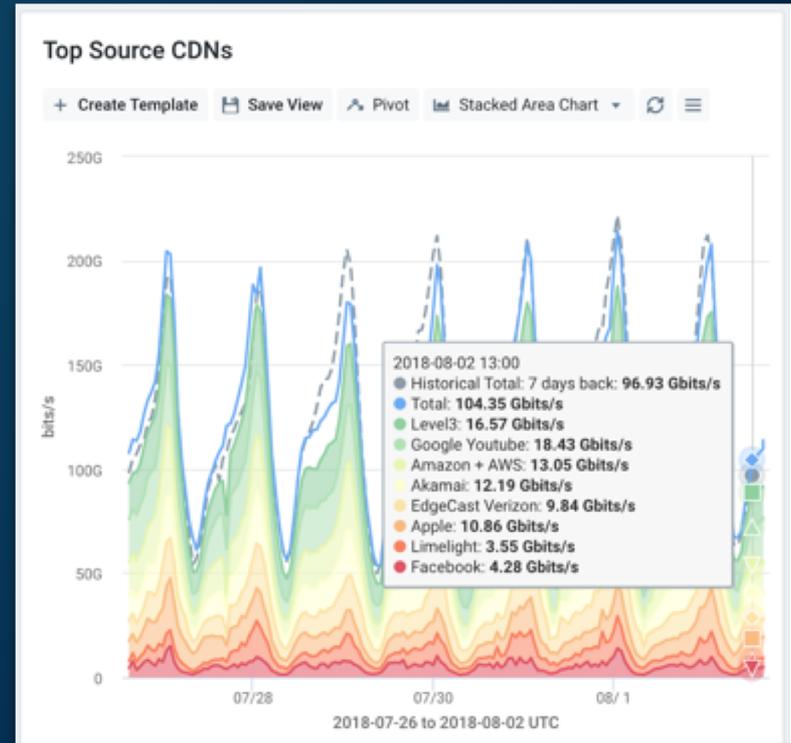


Each CDN delivers for multiple OTT providers



Subscriber and Service Utilization Analytics

- Real-time enrichment of network traffic data
- Adds context for business insight
 - Subscriber IDs (usernames, MAC, tier)
 - OTT service names
 - Originating CDN
- Understand network utilization per subscriber
- Or across subscriber segments
- Reduce customer service caseload
 - “Why am I over plan / quota?”
 - “Why is my connection slow?”
- Optimize traffic delivery from CDNs for cost and performance
- Create more profitable product pricing and packaging



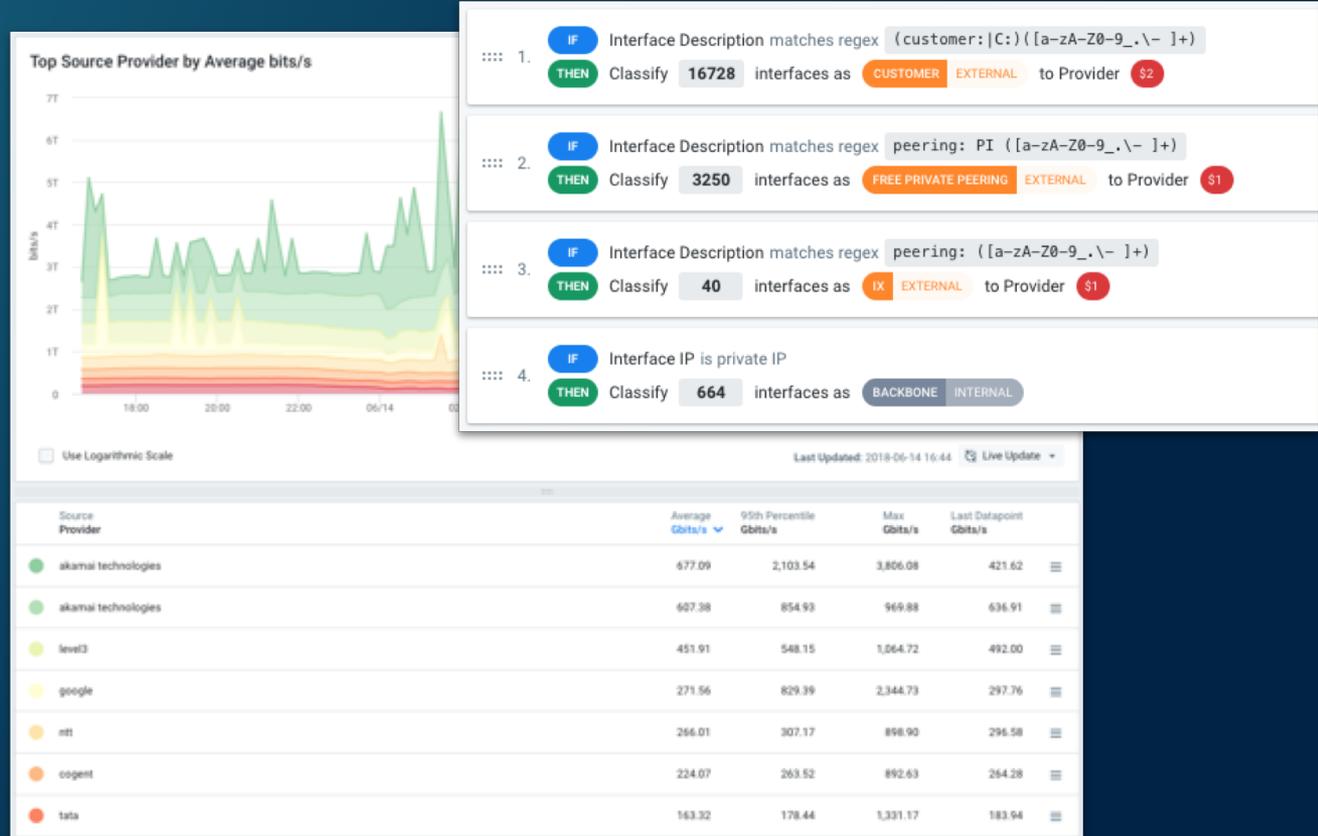
Interface Classifications

Filter your traffic based on interface or device roles, and the providers/customers behind

Devices

Accurately classify your flow data by identifying roles for your device interfaces. Is this an internal or an external interface? If this is external, is it for Peering, or Transit or for a Customer?

Interface Classification gives you tight controls for classification that help you filter to the exact view of the data you need.



Over-the-Top (OTT) traffic enrichment

- **Hard, but feasible**
 - OTT providers rely on owned infrastructure and CDNs
 - Combine Flows + DNS query data + Curated host patterns
 - Still done near real time at ingest.
 - A high cardinality / frequency flow tagging backend is required
- **Business impact**
 - **Identify** traffic or cache embedding opportunities
 - Additional, **end-to-end end-user support tool**





Misc Appendix



Kentik Ops

- Containerized microservice architecture
- Hybrid of private cloud in Equinix, + cloud proxies in AWS, Google
- In Equinix
- Our own provisioning system
- Linux + Docker + ZFS
- All nodes PXE
- MX + QFX network stack
- In-house backend for ingest, fusion, column store, streaming
- Go, Rust, C, C++
- Gigabits inbound via Internet, interconnection
- Unencrypted UDP and encrypted TCP

SecOps

ebay

tierpoint

IBM Cloud

GitHub

DreamHost

UBER

neustar

DigitalOcean

Bank of America

Once in Use,
Kentik Spreads

BizOps

TATA

gtt

Limelight
NETWORKS

cogent

CenturyLink

Bell

T

DevOps

credit karma

GitHub

IBM Cloud

Expedia

CISCO

AppNexus

MediaMath

DigitalOcean

twitch

IMPERVA

GoDaddy

CBS Interactive



What's New



New Functionality

Cloud Infrastructure

- AWS, GCP Flow Logs
- Auto-provisioned host monitoring
- CDN logs

Service Provider

- Subscriber Intelligence
- CDN Tracking
- OTT Tracking
- My Kentik Portal

Automation

- ServiceNow
- FlowSpec
- OpsGenie

New Functionality: Behind the Scenes

Enrichment - Scale

- Tens of millions of tags
- Updated in < second
- Millions at a time

Enrichment - Integration

- Kubernetes pod/service
- AWS tagging
- NSX tagging

User Interface

- Content library
- Universal search
- Guided dashboards
- Linked dashboards
- Expanded query
- Geo visualizations
- Performance bracketing
- Raw flow viewer



What's Ahead

kentic.com/nfd19



Negative Road Map

- Up / down pingthings
- Complete config parsing / “what if”
- App stack instrumentation (APM)
- Generic logging platform
- Generic BI
- IGP analytics/forensics
- Storage monitoring
- Wireless monitoring
- DB analytics
- SIEM
- Endpoint patch / version management
- Help desk / ticketing

What's Next: Q4 '18, 1H '19

- Azure logs, tap
- Metrics, Streaming Telemetry
- Segmentation, policy intent
- Wider CDN, load balancer log integration
- Turnkey orchestration integration
- White label cloud visibility platform
- Edge compute/deployed Kentik

What's Next: 2H 2019

- Traffic-based synthetic measurement
- INT / p4 integration