NetPanel: Traffic Measurement of Exchange Online Service

Yu Chen[1], Liqun Li[2], Yu Kang[2], Boyang Zheng[1], Yehan Wang[1], More Zhou[1], Yuchao Dai[1], Zhenguo Yang[1], Brad Rutkowski[3], Jeff Mealiffe[3], Qingwei Lin[2]

Microsoft 365, *China*[1] *Microsoft Research*, *China*[2] *Microsoft* 365, *USA*[3]

Exchange Online: Cloud based messaging platform

Component: performs a specific function as an entire or part of a process

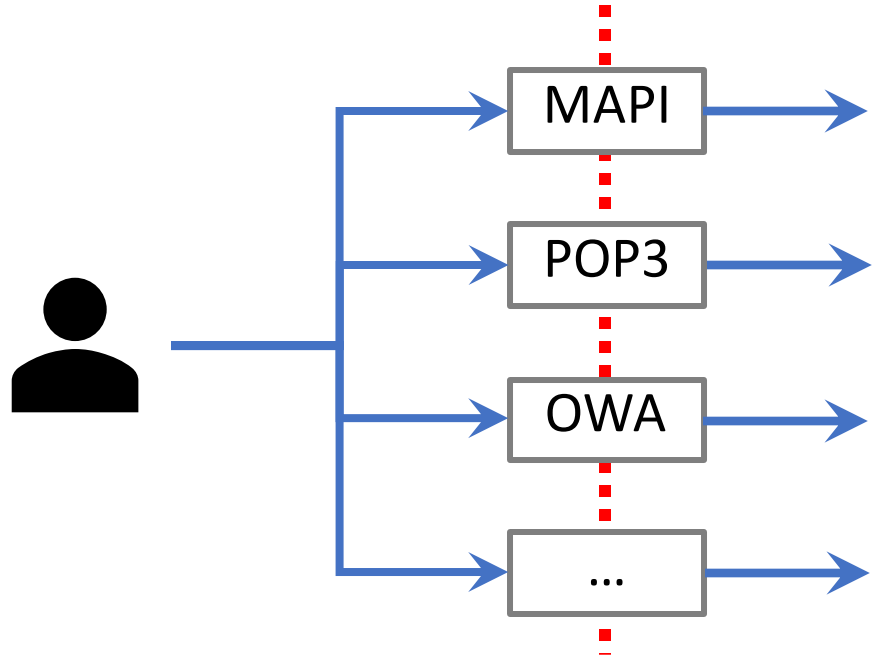

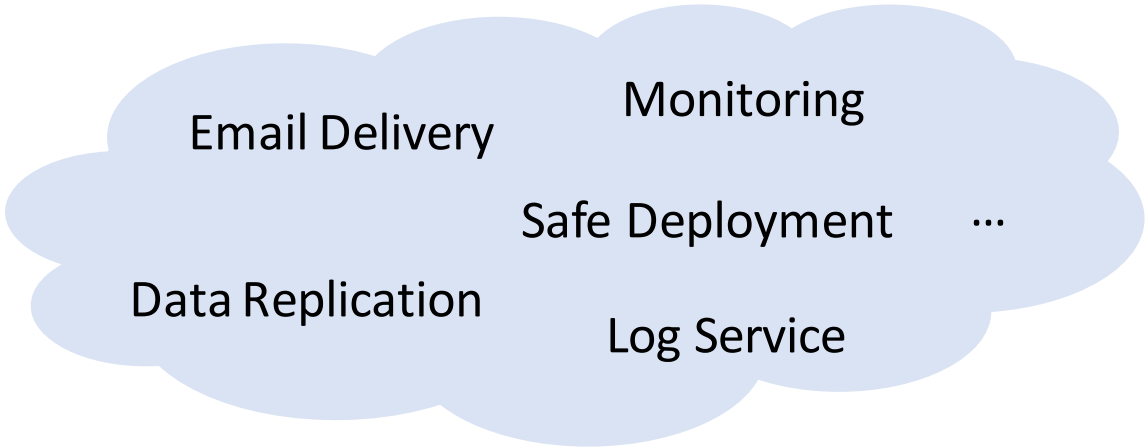

$$$ Hundreds of millions of dollars / year

Request Flow

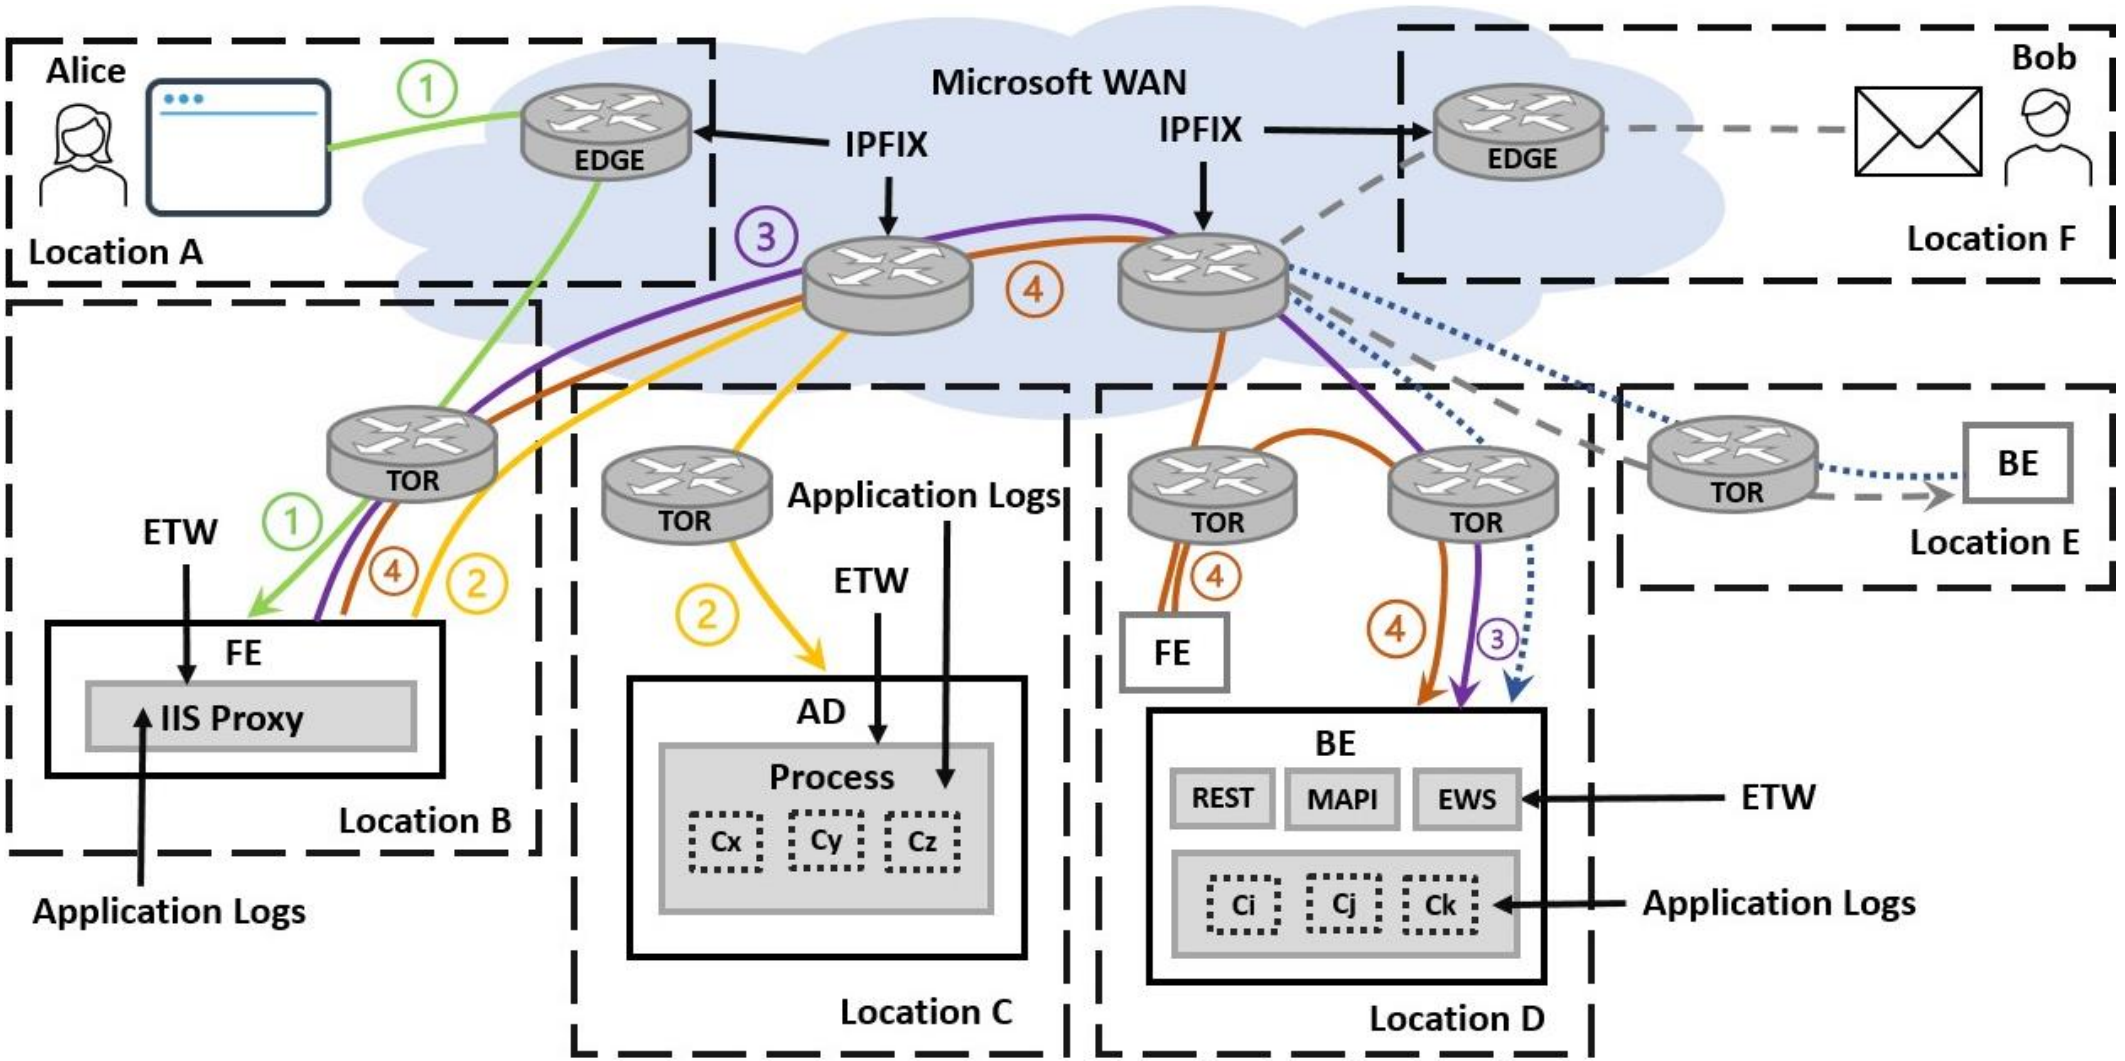

NetPanel Design Considerations

- Component level measurement
  - A single engineer team
  - Could be part of a process
- GB Level daily data size
  - Consecutive analysis over weeks of data
  - Originally, IPFIX more than 10TB
- Limited production overhead
  - High Service Level Agreements
  - Restricted resource consumption

NetPanel Overview

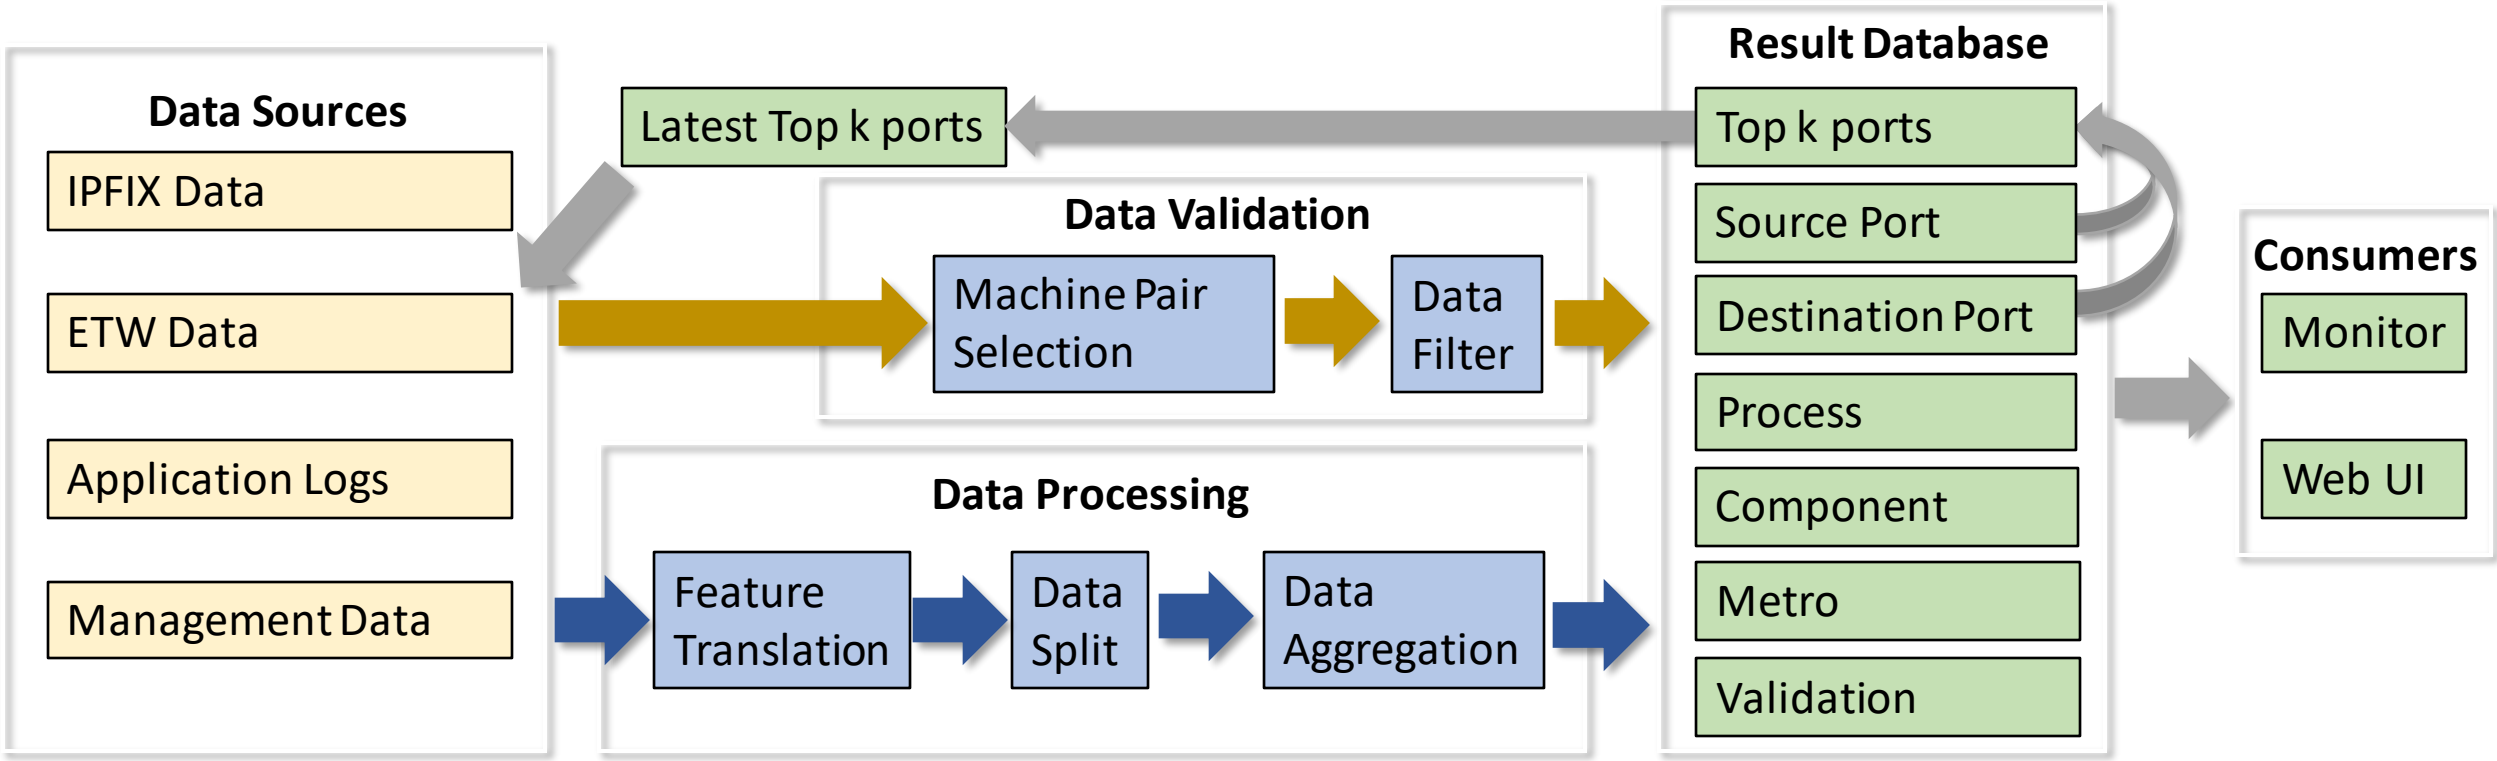

Data Processing

- Feature translation
  - IP Address -> ServerRole (AD, BE, FE, etc.), Location

- Data split and aggregation
  - Aggregate on source port and destination port separately
  - Top-k port to filter ephemeral ports

Data Validation

- Recovered traffic for a single pair of machines in one day from IPFIX matches the traffic from ETW

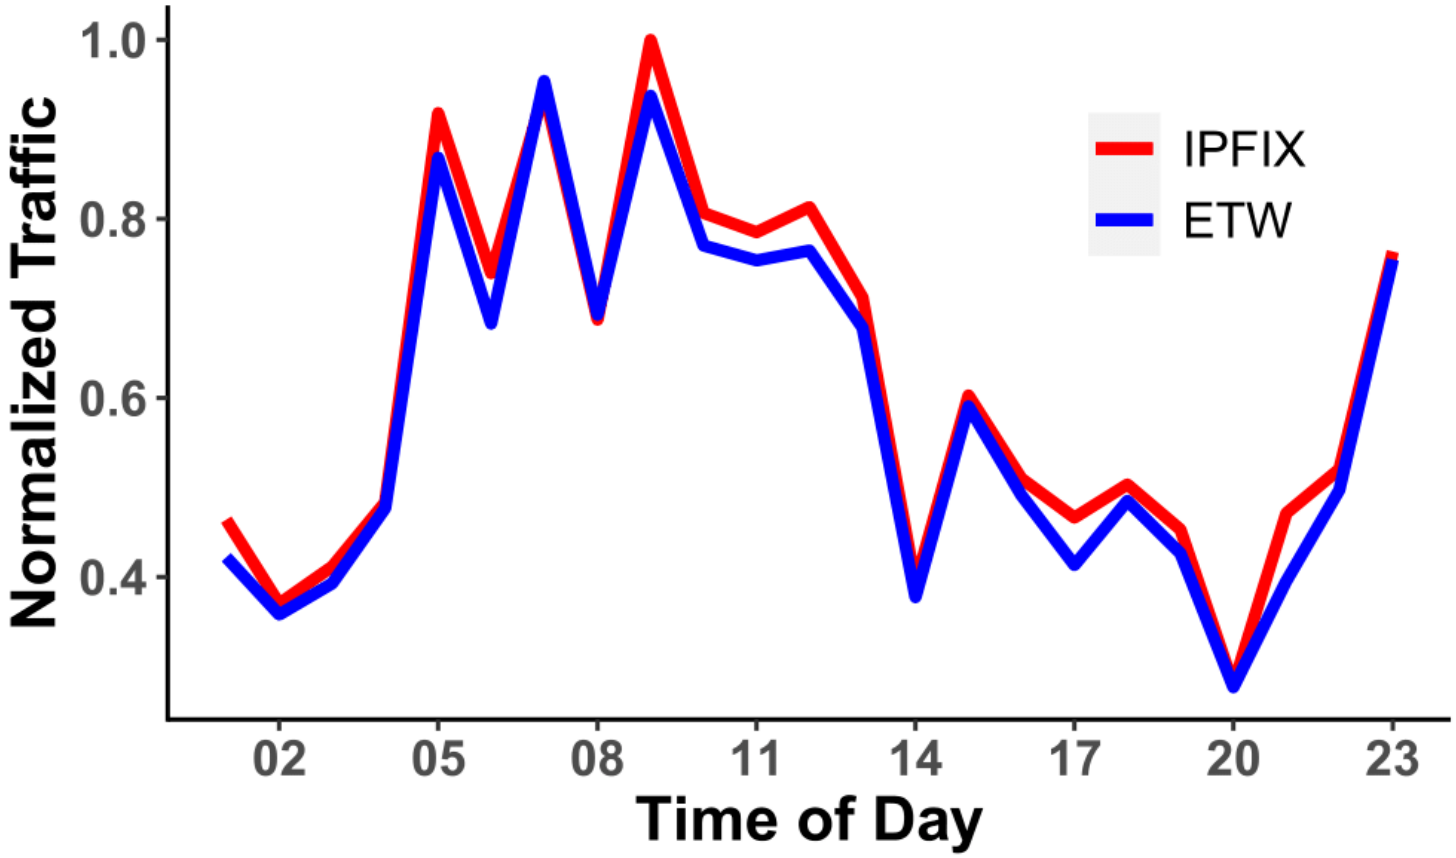

$$IPFIXBytes = \frac{(PacketSize + HeaderSize) * PacketNumber}{SamplingRate}$$

Production Insights

- Several ports dominate the overall traffic

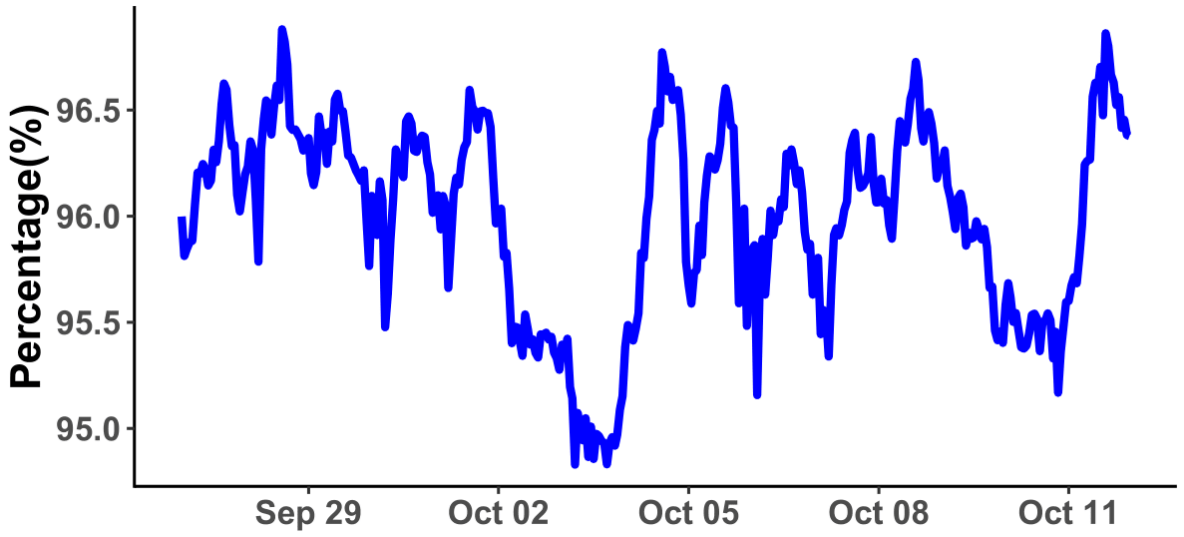

- Several components dominate the traffic of a top port

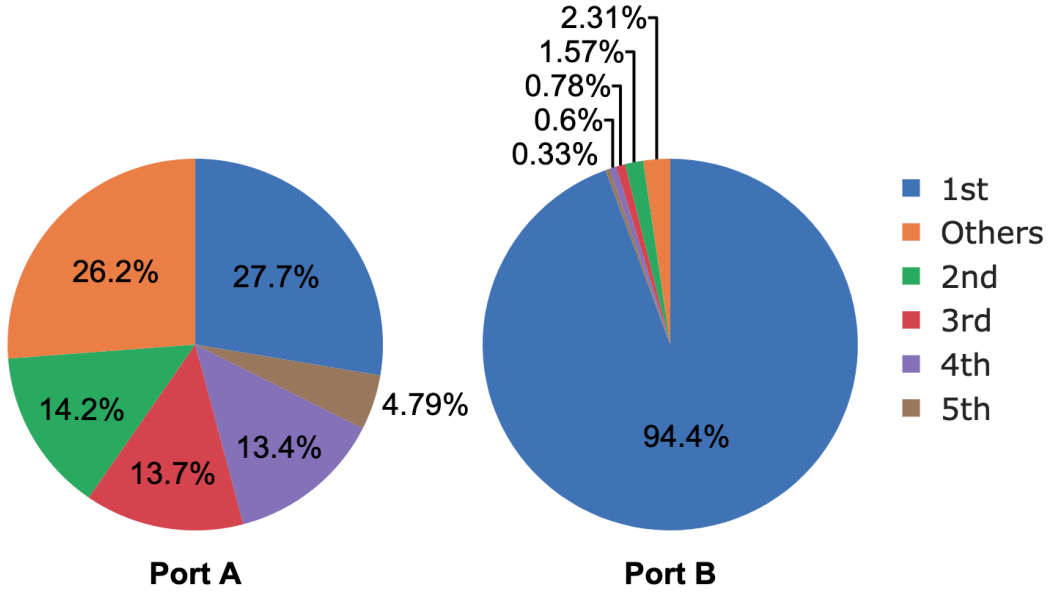

Case Study 1

- Service Traffic Optimization
  - Unexpected longhaul traffic from BE to FE
  - Internal service uses an endpoint without routing optimization

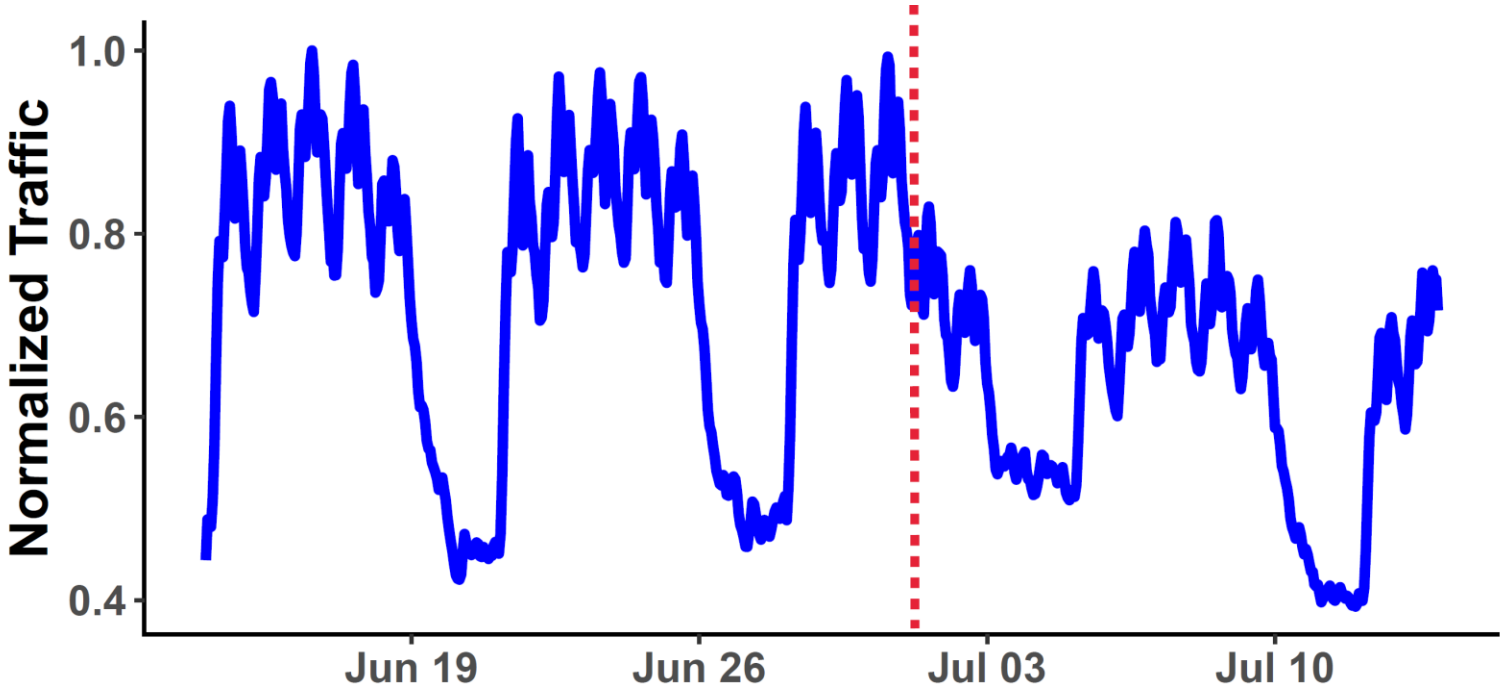

Case Study 2

- Legacy Traffic Discovery
  - Hard to detect if it does not break anything
  - Unexpected large traffic to port C

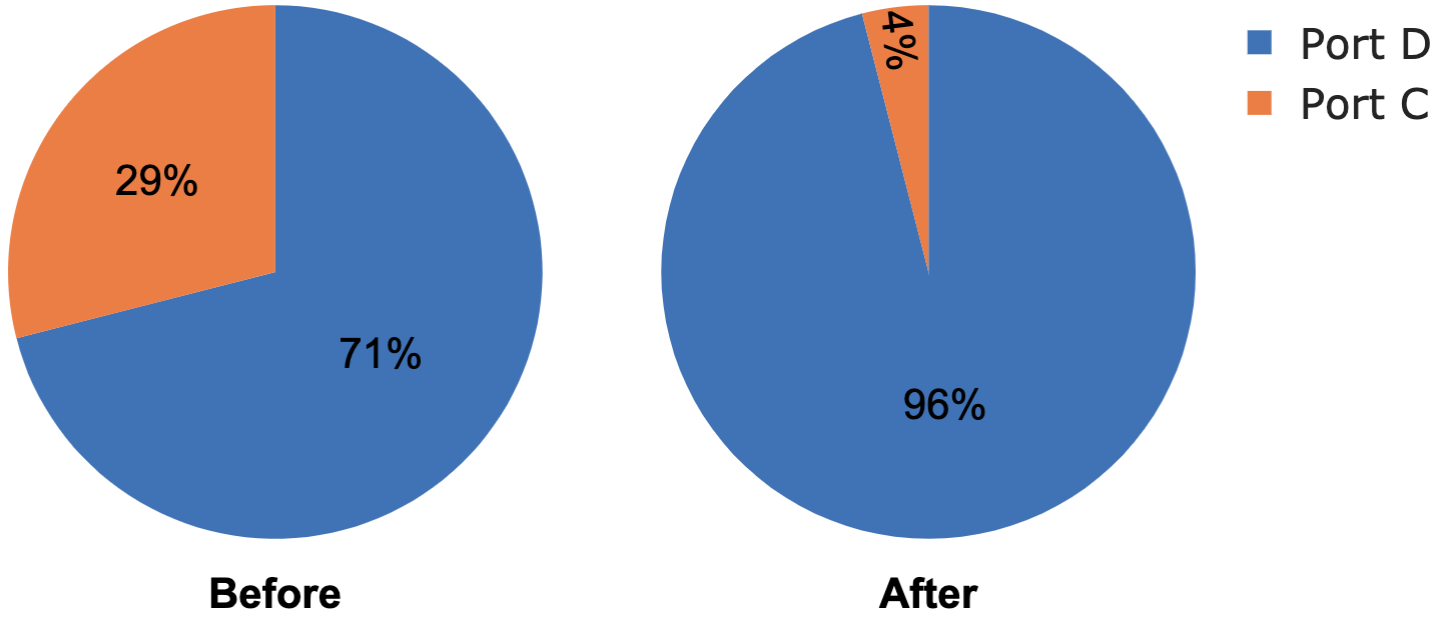

Case Study 3

- Anomaly Traffic Burst Detection
  - Accidently change the compression algorithm to older version
  - Quadruple peak traffic

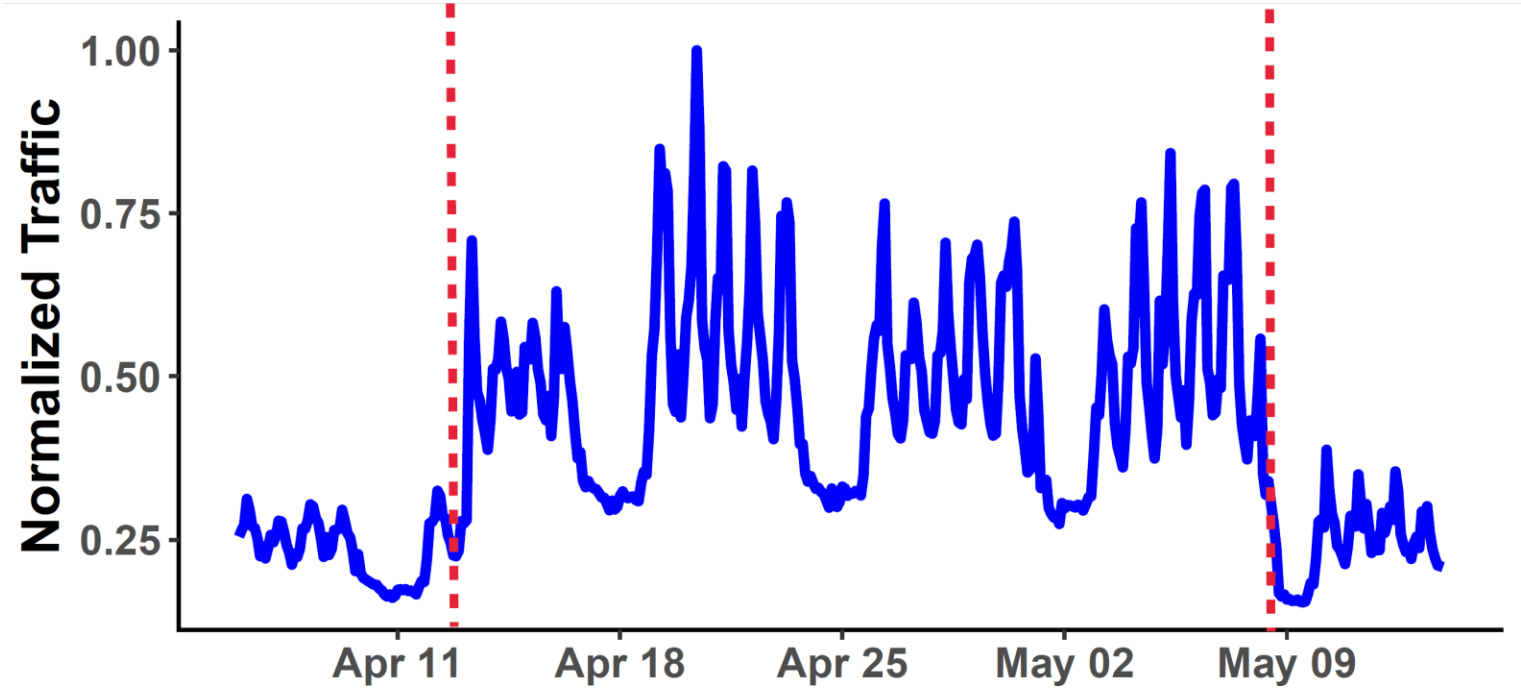

Case Study 4

- Wide-Area-Network (WAN) Feature Validation
  - New Device doesn't support existing feature well
  - Traffic priority out of expectation

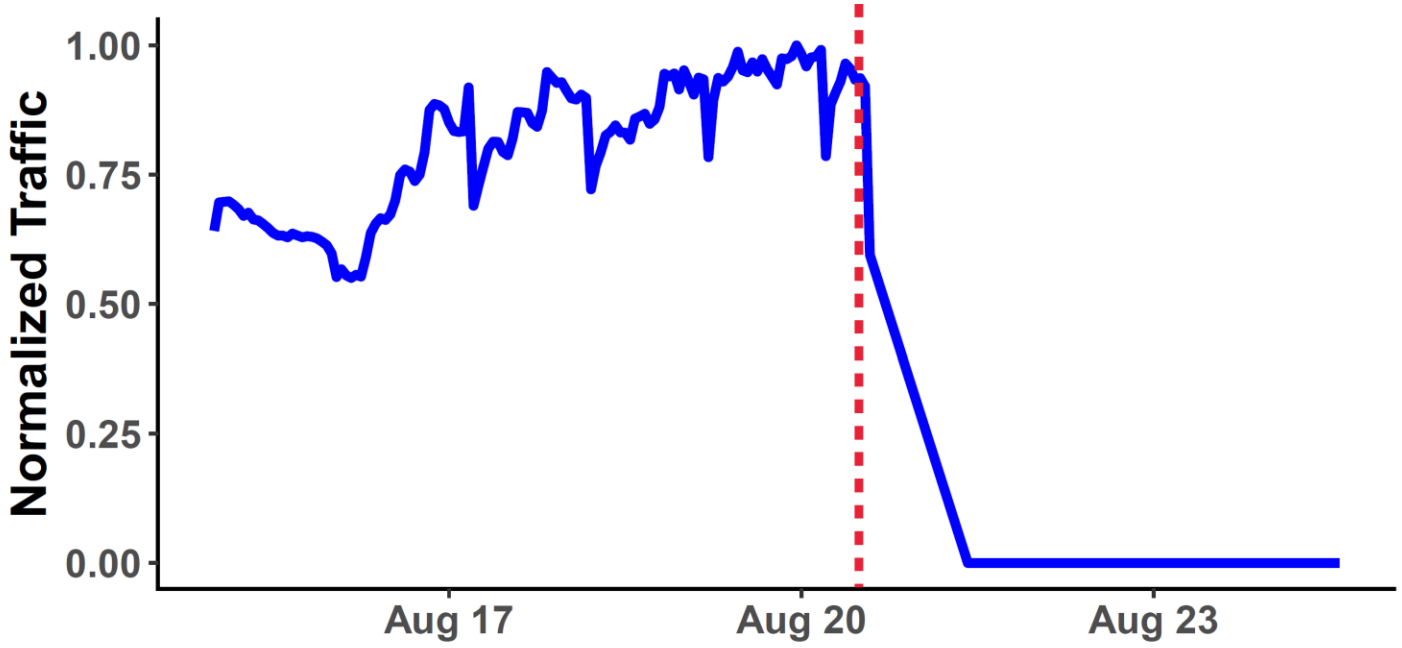

Evaluation

- In-production Overhead
  - Less than 1% rise in CPU and Disk IO
  - Saved 99.1% CPU and 99.7% Disk IO compared with raw ETW
- Outside-production Overhead
  - Tens of seconds to query 60 days of data

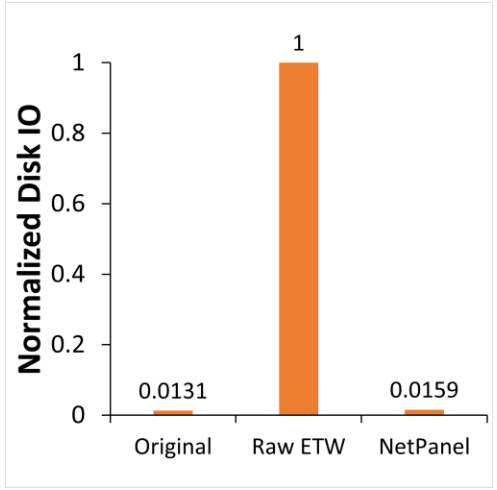

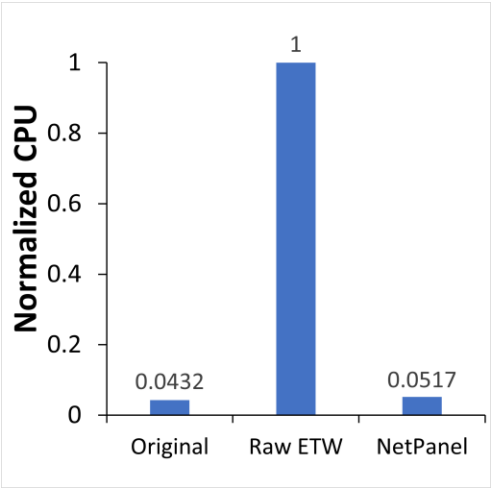

Data	Size Ratio	Daily Calculation
IPFIX	0.00361%	1.1 hours
ETW	0.00076%	2.5 hours
Application Log	0.00003%	6.8 hours

Conclusion

- Component level measurement is needed to drive the cost saving effort for a cloud application developed by a large number of teams.
- Heavy hitters are stable in EXO.
- NetPanel achieve component-level measurement leveraging IPFIX, ETW and application log.
- NetPanel has been in EXO production environment for more than 1 year and has helped us saved tens of millions of dollars.

Thank you!

Contacts: m365netcogs@microsoft.com