Overview

PPE | Programmable Packet Engine

The mantis Programmable Packet Engine (PPE) deploys in to, or alongside of, a production network and ingests raw network data packets. The packet traffic can be sourced from a network bypass TAP, a packet broker, or a SPAN/Mirror port. Internally, Containerized Virtual Engines (CVEs) target individual network protocols, packet attributes, and time series information to generate serialized metadata for consumption into a streaming data pipeline (i.e. Kafka, Kinesis). This level of "source-data-control" significantly reduces the volume of data being sent to follow-on processors, and provides a scalable approach to adjunct descriptive and predictive network analytics.

The PPE is a real-time, stream processing appliance that can be used with data-at-rest and data-in-motion system architectures and delivers real-time metadata streaming for multiple 1G and 10G network segments.

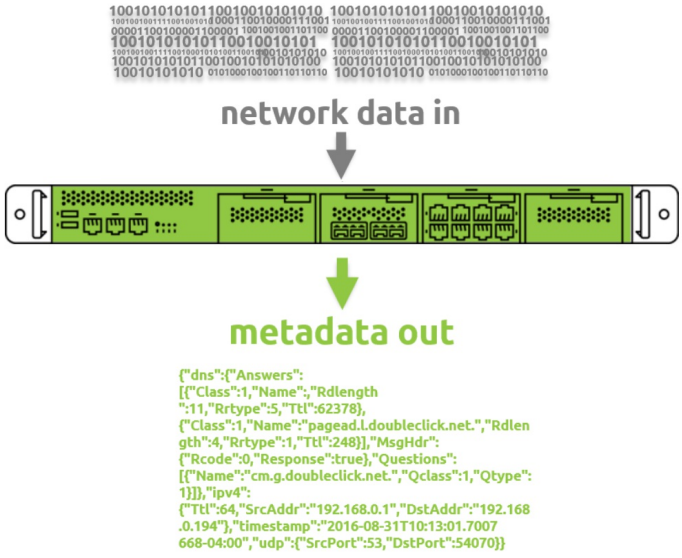

How it works

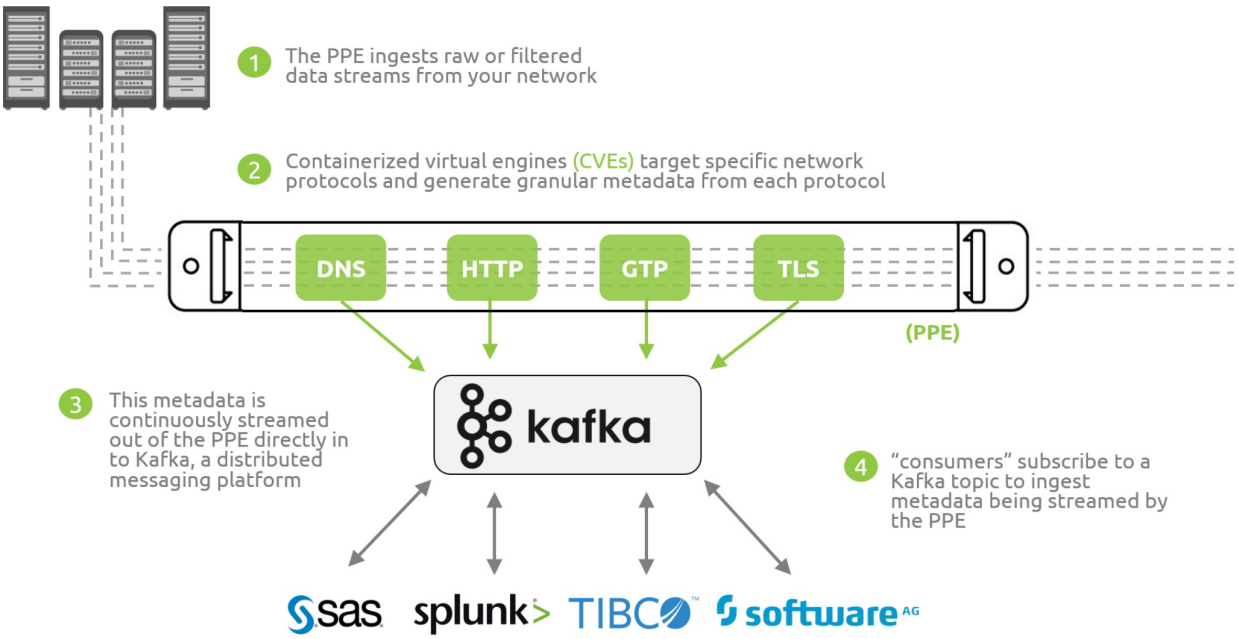

Why rely on network metadata?

The volume, variety, and velocity of data in contemporary IT infrastructures make network visibility both a **BIG DATA** and **FAST DATA** problem. Full packet capture from network traffic is a useful, though incomplete answer for network professionals. Data is simply too massive in scale, and complex in nature for organizations to copy every single packet to disk. OPEX and CAPEX costs involved are enormous- especially given the projected data growth rates of network traffic over the next ten years.

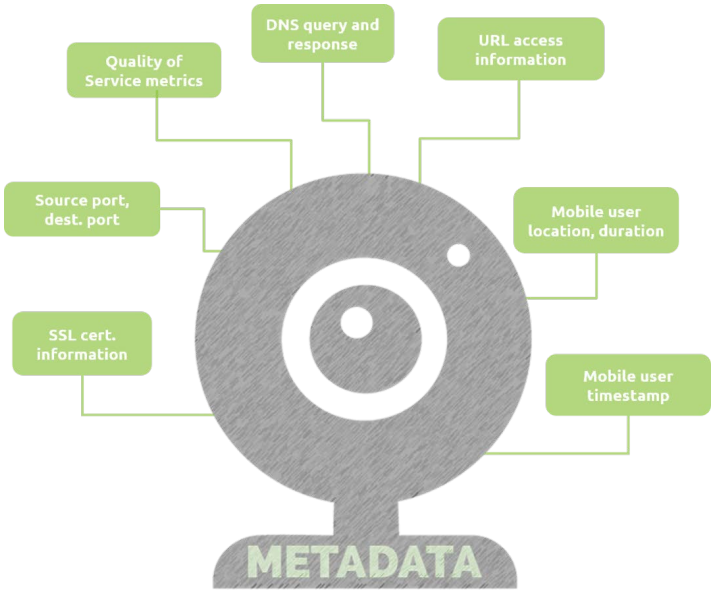

Wire-speed streaming analytics requires **intelligent, targeted extraction of metadata** from the underlying network protocols in order to analyze information found in today's networks. Ingesting real-time, protocol-specific metadata from the PPE allows downstream analytic applications to receive the most efficient view of network traffic, where all unwanted "network noise" has been removed and all permitted traffic is streamed in a very light-weight, but highly contextual form. Real time metadata streaming focuses on addressing both the VOLUME and VELOCITY challenges found in contemporary IT networks.

---

mantis METADATA

- ✓ granular, protocol-specific
- ✓ continuously streams in real-time
- ✓ streams can be accessed by any number of applications
- ✓ True FAST DATA pipeline

HTTP, DNS, GTP, TLS, DNP-3, Modbus, REGEX, ENTROPY (CVEs available today)

"Ingest-ready" metadata formats:

The generated metadata can be packaged in a variety of selectable industry standard formats (JSON, PROTOBUF, AVRO, OTHER). This metadata is forwarded into a KAFKA message queue for immediate consumption by consumer/subscriber nodes in BIG DATA/FAST DATA architectures. This design allows the PPE to immediately be integrated into a variety of predictive analytic and stream processing pipelines available today.

JSON protobuf Protocol Buffers AVRO

PPE appliance: hardware specs

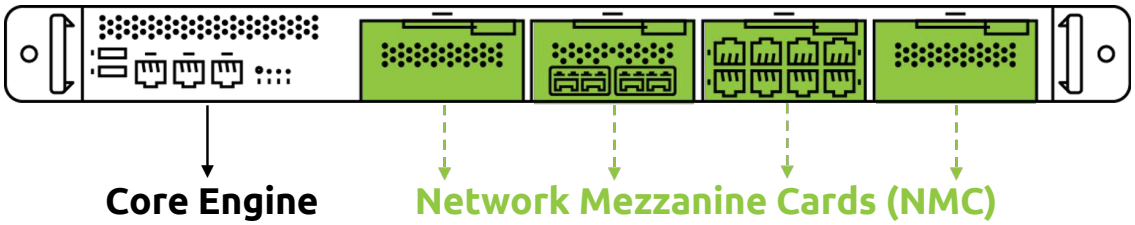

Core engine

The core engine is the main component of the PPE. It is a compact, 1U box with processing power to generate metadata from multiple 1 or 10G network streams

	PPE core engine
Processor	(2) Intel® Xeon® E5-2600 v3 / v4 Socket
Ports	(2) RJ-45 ports for management
Network Connectivity	(4) NMC slots (max 1G ports=32, max 10G ports=16)
Dimensions	17.24" x1.732" x24.61"
Power type	AC, redundant or DC, redundant
Watts	650W
Operating Temperature	0 ~ 40 °C (32 ~ 104 °F)

Network Mezzanine Cards (NMC)

Each core engine comes with 4 **network mezzanine card (NMC)** slots- which are used for ingesting various network streams and types. Customers choose from a variety of NMC cards, for both 1 and 10G. All cards include integrated bypass.

Part number	Port type	Speed	Port count
PPE-NMC-4C	RJ-45	up to 1G	(4)
PPE-NMC-8C	RJ-45	up to 1G	(8)
PPE-NMC-2SRF	LC (SR)	up to 10G	(2)
PPE-NMC-2LRF	LC (LR)	up to 10G	(2)
PPE-NMC-4SRF	LC (SR)	up to 10G	(4)
PPE-NMC-4LRF	LC (LR)	up to 10G	(4)