ExtraHop Briefing Notes

Author: Ronnie Beggs.

ExtraHop is a breath of fresh air in the crowded world of IT Operations Analytics (ITOA). ExtraHop performs real-time stream analysis of the packets that carry data across a network, transforming those packets into structured wire data (as it is described by ExtraHop). Packets, or network data, have been used widely for decades by network engineers to identify protocol and communications issues, yet the network has remained relatively untapped as a source of operational intelligence for mission-critical applications.

Most approaches to operational intelligence rely on the software applications to generate meaningful management data, either through instrumented APIs for performance metrics, or as the log data generated directly by the application. According to ExtraHop, this drives multiple siloes of analytics tools, driven by data type, and that the agent-based and instrumented API models for data collection require systems to be hard-wired in advance for specific metrics and data.

With wire data, ExtraHop believes it has a better way. Transforming raw packet data at volumes of hundreds of terabytes per day into meaningful analytics, and across a wide range of industries and use cases, requires a highly scalable and flexible platform. The rewards though are obvious - wire data holds real-time insights into all network communication between all applications, devices and services.

The task of decrypting, decoding and assembling application-level information from protocol packet data should not be underestimated. It requires significant processing power, the right architecture, and some clever software. If you’re familiar with the world of protocol analyzers and probe-based network fault finding in the telecoms sector, this can be a deeply technical area, requiring significant network and protocol expertise. The ExtraHop proposition is that life does not have to be this way, it’s about how the information is presented and the facilities available to the user to analyse and understand the data. This is an interesting aspect of ExtraHop, the use of UI tools to simplify the underlying complexity from a user’s perspective.

The core of the ExtraHop platform is the Discover Appliance, available as a physical, virtual, or cloud appliance. The physical appliance is a 1U or 2U rack mounted unit that is installed in the network data center, or a small form factor unit for remote offices. To deliver analysis, each Discover Appliance requires a copy of the data flowing across the network, which can be achieved by cabling (physically or logically) the Ethernet capture ports on the appliance to taps or port mirrors on the network equipment. This requires data center access and specialist network skills, and can be non-trivial in some environments, but nothing that is out of ordinary. Once the appliance receives the copy of network traffic, it reassembles the packets into complete transactions, flows and sessions with analysis of all the data across all of protocol layers from the data layer (L2, for example Ethernet), through the network layer (L3, for example IP), up to the application layer (L7). No agents for data collection are required, nor instrumentation of software systems for metric collection.

Copyright © Bloor Research 2016

The physical Discover Appliance can process up to 40 Gbps, monitor up to 5,500 servers, and process in the order of 432 Terabytes of data per day. Multiple Discover Appliances can be deployed for larger networks, with overall control administered from a single Command Appliance. Each Appliance is configured to extract the protocols and payloads of interest. Local data storage is based on fast file-based storage for performance, with the capacity to retain a minimum of 30 days of monitoring data and metadata. The Discover Appliance can also connected to external NAS drives.

The ExtraHop user base includes the network engineers, with UI tools for deep dive transaction analysis, but moving towards operational intelligence and operational users who have different reporting and analytics requirements. Users access reports directly on the Discover Appliance (or through the Command Appliance if installed) over HTTPS and have access to a range of customizable dashboards, alerts and views for both network engineers and operational users. There is also an Explore Appliance, built on ElasticSearch, for longer term data storage and,analysis. The Explore Appliance is populated by transaction and flow records from each Discover Appliance, with data and analytics integrated in the same UI. The Open Data Stream (ODS) interface offers a mechanism to stream data to other systems, including Elasticsearch, MongoDB, Kafka (and Hadoop), and Splunk.

IT Operations is ExtraHop's primary market but with presence in the Internet of Things (IoT) and Healthcare. For example, one healthcare provider is using ExtraHop to analyze HL7 messages (a healthcare specific protocol for the transfer of clinical data between applications and healthcare providers), in order to identify and alert on unusual drug prescription activity in real-time. Competitors in the ITOA wire data market include Riverbed and Netscout, but with real-time stream analytics, ExtraHop believes it has a significant competitive advantage.

ViewPoint

ExtraHop is built on data networking and appliance expertise gleaned from the founders' previous experience at F5 Networks. We believe the core customer base in ITOA will remain buoyant as Extrahop offers a compelling solution on price, scale and usability. The expansion of their market footprint has the capability to be disruptive and will be a development we will follow with interest. First, in the ITOA sector, vendors such as Splunk and AppDynamics may start to feel increasing competitive pressure from ExtraHop. Also, the potential for disruption exists in the wider Internet of Things market for monitoring and operational analytics, where passive monitoring of network traffic eliminates the need for agent instrumentation and complex integration. ExtraHop's experience in the appliance market—of the remote configuration and management of appliances with data synchronization—also offers the right experience and technology for success in the IoT market where edge analytics and appliances will soon be the norm.

Copyright © Bloor Research 2016