

Flow Analysis in Heterogeneous Cloud Scenarios

Jochen Kögel, Sebastian Meier, Stefan Oettl
Isarnet Software Solutions GmbH, Munich, Germany,
{jochen.koegel, sebastian.meier, stefan.oettl}@isarnet.de

Abstract—As enterprises move their private and public services to the cloud, traffic monitoring between clouds becomes crucial and challenging. Compared to on-prem resources, the cost structure in the cloud is very different and requires a good understanding of application topologies and resulting traffic patterns before, during, and after (partial) migration to the cloud. This requires traffic monitoring systems that provide a comprehensive view across enterprise WANs, on-prem services, and multiple cloud providers. We show from an industry-perspective how the well-established NetFlow/IPFIX data sources and Flow Logs available on a file basis from cloud providers can be combined in one solution that provides visibility as a single pane of glass to network operations and application owners. We introduce the IsarFlow Flow Logs Transcoder, its features and integration in Amazon Web Services (AWS) environments.

Index Terms—Flow Analysis, Cloud Infrastructure, Network Monitoring

I. INTRODUCTION

Flow Logs can be used to record and store information about network traffic in AWS VPC environments. As with NetFlow/IPFIX, traffic data is captured at flow level. The IsarFlow Flow Logs Transcoder converts traffic data from Flow Logs to standard-compliant IPFIX and forwards the IPFIX data stream to a configurable destination (Fig. 1). Monitoring of inter/intra-VPC traffic and traffic towards on-premises networks and other AWS regions is thus seamlessly integrated into our traffic monitoring solution IsarFlow [1].

Heterogeneous cloud environments are characterized by a mixed infrastructure of on-premise resources and virtual private clouds spread across various public providers (Fig. 2). In such scenarios, data traffic from clients to services as well as between various application servers and resources flows between cloud, on-prem infrastructure, office locations as well as the Internet.

Flow data for traffic analysis is not only available from, e.g., edge routers via NetFlow/IPFIX, but also from cloud providers such as AWS or Azure via Flow Logs. In order to provide




Fig. 1. Flow Logs Transcoder overview.

an overall understanding on traffic amount and characteristic, a comprehensive view incorporating all flow data sources is crucial. For example, to understand performance issues on certain application flows, it is required to investigate whether bottlenecks are in on-prem or cloud network functions. Furthermore, traffic cost in the cloud depends on properties, such as traffic direction and regions where services reside. Hence, an overall understanding especially for enterprises with global on-prem and cloud-deployments is essential. In the following, we will show how we integrate Flow Logs and NetFlow/IPFIX into IsarFlow to provide this comprehensive view.

II. FLOW LOGS TRANSCODING CONSIDERATIONS

Converting AWS Flow Logs needs certain understanding of AWS deployments, resulting monitoring data, and IPFIX semantics. This section details the corresponding requirements and considerations.

A. AWS Flow Logs sources

In the following, we focus on AWS Flow Logs and deployments, while Azure Flow Logs provide similar data.

Network telemetry data can be collected at various observation points in AWS environments:

a) *VPC Flow Logs*: VPC Flow Logs can be activated at the interface, subnet, or VPC level.

b) *Transit Gateway Flow Logs*: A Transit Gateway Transit Gateway (TGW) connects several VPCs and on-premise networks as a central hub. Flow Logs can be activated for individual TGW attachments or for any traffic traversing the TGW.

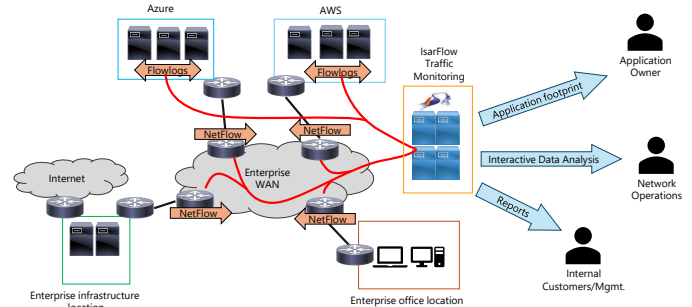


Fig. 2. Heterogeneous enterprise cloud scenario with traffic monitoring using NetFlow/IPFIX and Flow Logs as data source.

B. Deduplication with TGW

Traffic flows traversing a Flow Log observation point (network interface) are always captured in ingress and egress direction with respect to that observation point. Therefore, traffic flows may be captured multiple times along their path in the monitored AWS environment.

This is most obvious in scenarios where Flow Logs are activated for an entire Transit Gateway (Fig. 3): In this case, a traffic flow entering the TGW is accounted at the ingress TGW ENI (Elastic Network Interface) and at the egress ENI. The same field values are in both Flow Log records. In order to avoid duplication of Flow Log information, this information must be properly processed and understood or filtered for a given direction only. For this reason, the flow log transcoder provides direction filtering capabilities.

C. Technical Details for IPFIX Encoding

For some AWS Flow Logs data fields, there is no straightforward mapping to IANA-assigned¹ IPFIX information elements (IEs) or data types. This could be solved by using enterprise-specific IEs. However, we decided to apply conversions to IANA IEs to make the resulting IPFIX data more general and more consistent with established data analysis approaches.

a) *Interface-IDs*: In AWS VPC environments, interfaces are identified by interface IDs. Interface IDs are of type string and use the format prefix-ID (e.g., eni-afdf9f8c5db2a7ba2). Since IPFIX uses integers for ingress and egress interfaces, AWS interface IDs are automatically converted into integer values by the Flow Logs Transcoder. In order to reassign the integer to the original interface ID on the collector side, information on interface IDs, interface names and interface descriptions is exported via IPFIX Option Records. This data is automatically processed by IsarFlow and converted into configuration data. From a user's point of view, the AWS interface representation therefore does not differ from conventional router interfaces within IsarFlow.

b) *VPC-IDs*: In AWS VPC environments, VPCs are identified via VPC IDs. In IPFIX, there are no native fields for VPC IDs. In the Flow Logs Transcoder, the VPC ID is translated into a VRF field, as a VRF has very similar semantics compared to the VPC ID. Similarly to interface

¹<https://www.iana.org/assignments/ipfix/ipfix.xhtml>

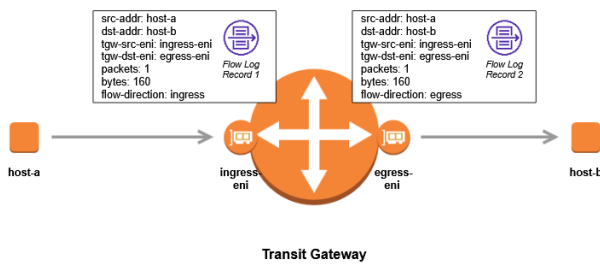


Fig. 3. Duplicate flows in TGW scenario.

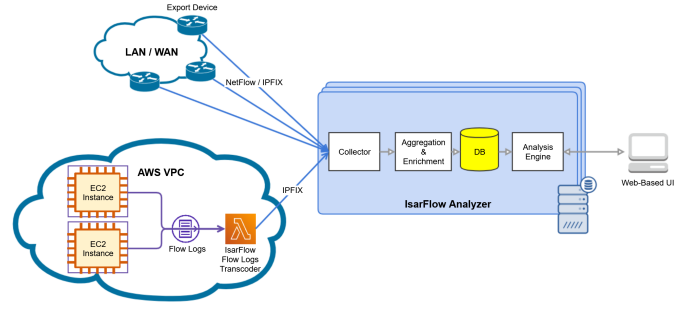


Fig. 4. Deployment for heterogeneous data source processing.

IDs, the AWS string representation of the VPC ID is also automatically converted into an integer value and transmitted via IPFIX. To be able to reassign the integer to the original VPC ID (VRF ID) on the collector side, information on VPC IDs, VPC names and VPC descriptions is exported via IPFIX Option Records. This data is automatically processed by IsarFlow and converted into configuration data. From a user's point of view, the AWS VPCs do not differ from conventional VRF instances within IsarFlow.

III. INTEGRATION

A. Flow Logs transcoding via AWS Lambdas

Flow Logs can be used to record and store information about network traffic in AWS VPC environments. As with NetFlow/IPFIX, traffic data is captured at the flow level. The IsarFlow Flow Logs Transcoder converts the traffic data from Flow Logs to standard-compliant IPFIX and forwards the IPFIX data stream to a configurable destination. The transcoder is realized as Lambda due to its simple deployment and no fixed cost compared to continuously running EC2 instances. Monitoring of inter/intra-VPC traffic and traffic towards on-premises networks and other AWS regions is thus seamlessly integrated into IsarFlow (Fig. 4).

B. Deployment considerations

1) *Transcoder Deployment*: Transcoder deployment takes place from the Serverless Application Repository into a VPC environment. The installation is mostly automated. Flow Logs can be stored in different target technologies (so-called data sinks). The Flow Logs Transcoder currently supports CloudWatch Logs and S3 Buckets. Since data processing depends on the target technology, a tailored serverless application is offered for each supported data sink.

For data processing, the Flow Logs Transcoder requires access to the following APIs, which must be reachable via corresponding VPC endpoints:

a) *EC2 API*: The EC2 API is used for the following tasks: querying the Flow Logs record format at CloudWatchLogs, resolving interface IDs to interface names, resolving VPC IDs to VPC names.

b) *S3 API*: The S3 API is only required for the Flow Logs-to-IPFIX-S3 application. The API is used to read the Flow Logs data stored in the S3 bucket.

