May 14, 2018

AUTOMATIC MEASUREMENTS AND ANALYTICS IN CLOUD BASED MEDIA APPLICATIONS

Alessandro Duminuco

Peter Bosch

Andre Surcouf

Jeffrey Napper

AUTOMATIC MEASUREMENTS AND ANALYTICS IN CLOUD BASED MEDIA APPLICATIONS

AUTHORS:
Alessandro Duminuco
Peter Bosch
Andre Surcouf
Jeffrey Napper

ABSTRACT

Techniques are provided to measure and analyze the "health" of a cloud-based professional media application. In particular, described are (1) a (media) probing infrastructure, based on a mix of virtual and physical measurement probes or collectors; and (2) a set of techniques that allow an orchestration system to control the probing infrastructure, to deploy and/or configure measurement probes, and/or to divert or mirror traffic to be "measured" through such probes. This approach not only uses virtual or physical probes, but also the automatic and policy-driven orchestration of such probing infrastructure and leverages Internet Protocol (IP) probing and measurement mechanisms for the media space.

DETAILED DESCRIPTION

The professional media industry is virtualizing and cloudifying their assets. The industry is changing from bespoke (network and Audio/Video (A/V)) elements and physical cabling between elements and entities to an infrastructure where physical devices such as cameras and displays are connected to virtualized, cloudified appliances. Especially during the transition between the old and new, professional media networks are comprised of a mixture of virtual and physical assets.

One of the key areas that need to be addressed in professional media systems is proper media "signal" propagation: that is, how a professional media system can guarantee end-to-end delays, zero signal loss, and minimized end-to-end jitter in the signal. Many professional media workflows are comprised of (service) chains of appliances and detecting errors in a physical deployment is already daunting. One needs to deploy probes in the chain to understand where the "signal" is getting corrupted, delayed, or otherwise

treated incorrectly, and there exists a whole host of physical measurement devices to detect how well appliances are treating the "signal".

When appliances are virtualized/cloudified and the "signal" is carried in Ethernet frames and sent as Internet Protocol (IP) packets, detecting flaws in the "signal" is more daunting. Deploying physical measurement points is harder because (a) the appliance is hosted someplace in a data center, connected via some real or even virtual switch or switches to other virtual appliances; and (b) the application topology is dynamic and controlled by an orchestration system which may place and connect components differently at each deployment or even modify topology at runtime.

Accordingly, a set of techniques is described herein which allows an orchestration system to control a (media) probing infrastructure based on a mix of virtual and physical measurement probes. These techniques may involve deploying and/or configuring measurement probes in order to divert or mirror traffic to be "measured" through such probes.

A mix of virtual and physical measurement probes may be used, including traditional physical media specific probes, network device based probes, and virtual probes. Traditional physical media specific probes (e.g., IP-based media probes) are currently in use. Their configuration may be automated to be used in virtualized media deployment, potentially as a service. Network device based probes (i.e., specially instrumented network switches or physical Network Interface Cards (NICs)) may (a) perform media-related measurements on the traffic traversing their ports and send such measurements towards collecting entities for further analysis (e.g., media-specific netflow), or (b) mirror (relevant) traffic to external measurement device.

Virtual probes are functions deployed in a virtual/cloud environment in the form of virtual machine, containers, or server-less functions. Similar to the physical devices, such functions are able to analyze the IP-based media stream (also referred to as "the media signal"), to perform media relevant measurements on the signal where possible, and to expose the resulting data to collecting entities. A virtual probe may be "in-line" (i.e., directly traversed by the media signal) or external, in which case the virtual probes can tap the network traffic or receive a copy of such traffic with the help of other media functions

implementing mirroring capabilities. In case of multicast media delivery, a virtual probe can receive the media signal by simply joining the multicast distribution.

Given the time sensitive nature of media flows, the probing infrastructure must support accurate timestamping to provide relevant media analytics. As such, both physical and virtual probes may make use of highly accurate timing whenever provided by the underlying NICs (e.g., Precision Time Protocol (PTP) support). Since it receives the actual signal, a virtual probe may also perform media level analysis such as, e.g., black frame detection, etc.

Virtual cloudified applications are typically deployed by means of one or many orchestration systems. As described herein, the functionalities of such cloud orchestration systems may be extended to control the deployment and configuration of the probing infrastructure. In one example, an orchestration system takes as input a high-level description of a cloudified (media) distributed application (referred to herein as a "solution model") and performs compilation and deployment. The compilation consists of applying a set of (compilation) policies to translate the solution model into an actual deployment recipe for application (referred to herein as a "solution descriptor"). The deployment consists of deploying and configuring the distributed application on (possibly many) cloud systems by following the "compiled" solution descriptor.

Moreover, a compilation policy is described herein which may perform several actions. For example, the policy may analyze the solution model and identify the function chains traversed by media signals (i.e., the media pipelines). Additionally, the points in the media pipelines where measurement probes are needed may be identified. This decision can be driven by policy configurations (e.g., probes might be requested at all stages of the pipeline, only when the signal is carried over the network, or only when the signal is in particular formats). Identifying these points can result from the application of a dedicated policy.

The best suitable probe type for each probing point may also be identified. Because physical probes might be available only on specific hosts or specific switch ports, and virtual probes might be capable of analyzing only certain media formats, the user may express preference over the probes to be used. The policy takes all this into consideration to make the probe selection.

For each probe that needs to be installed, the solution model may be augmented with the needed configuration according to the probe type. This may involve inserting the inline virtual probe, ensuring it is traversed by the signal, and configuring it with the appropriate measurement and shipment instructions. Alternatively, this may involve inserting a virtual probe and ensuring it is configured to "tap" the signal, receive a signal mirror, or join a multicast tree if media is distributed over multicast. In another example, the appropriate network switch ports may be configured to enable the measurements or mirroring and to ship them where needed. In yet another example, the Software Defined Networking (SDN) controller may be instructed to funnel or mirror the (appropriate) media traffic to the physical probes that need to be used.

Figure 1 below illustrates a topology including a sample media application.

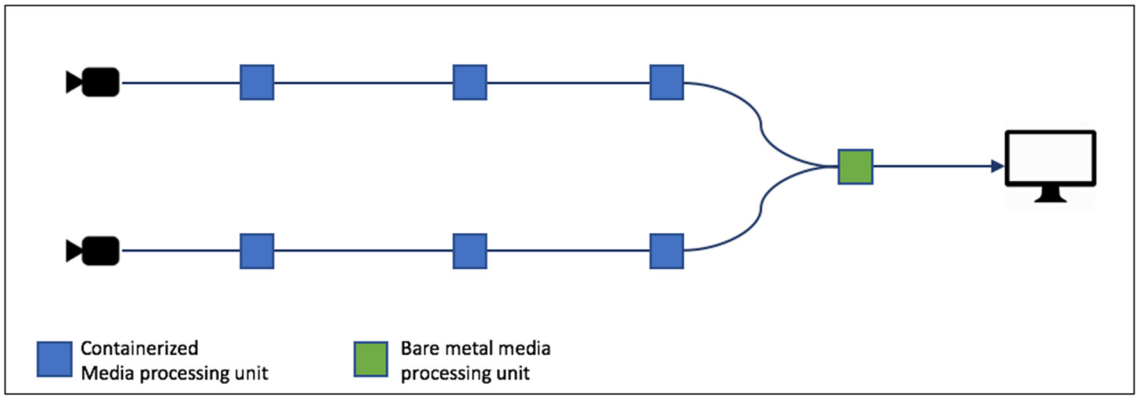

Figure 1: Sample media application – solution model

This solution involves two media sources (e.g., two cameras) feeding two media signals into two parallel media function chains, each composed of three container-based media processing units performing various functions such as logo-insertion, upscaling, etc. The outputs of the two chains are then sent to a bare-metal media processing unit, where the two signals are combined (e.g., with a multi-viewer or a composer). Finally, the output of the composer is sent to a media sink (depicted as a display).

In an orchestration system, the sample application is described by an abstract description (referred to herein as a "solution model"), which is analyzed and compiled to produce a solution descriptor. In the application, the compilation consists of a set of policies that define application details left undetermined in the model. For example, a compilation policy may determine the transport to be used among the various components,

while another compilation policy may decide which log-level to enable in the various components according to the deployment needed. Yet another compilation policy may determine the mapping between processing units and data center (physical) hosts (in case it has not already been defined in the solution model). Figure 2 below illustrates an example of the result of the compilation, where each processing unit is assigned to a specific physical host.

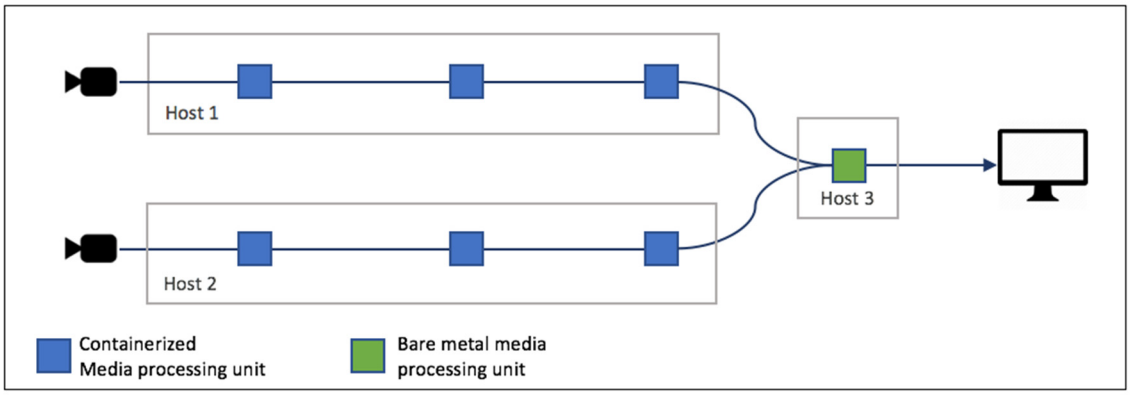

Figure 2: Sample media application - compiled solution model

A dedicated policy (referred to herein as a "probe policy") enriches the solution with the appropriate probing infrastructure and configures the infrastructure accordingly. Figure 3 below illustrates an example of a result of such a policy.

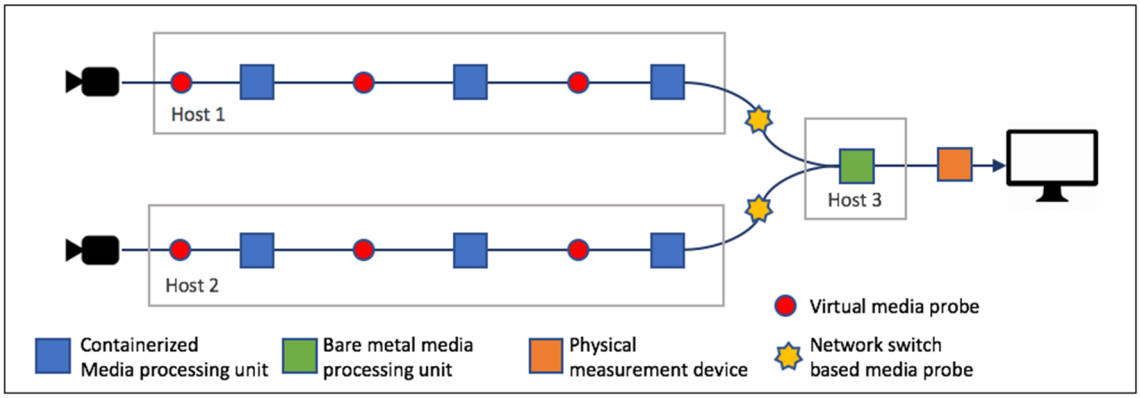

Figure 3: Sample media application - result of analytics policy

In this instantiation of the policy result, a mix of three types of probes has been added and configured in the infrastructure. First, a set of virtual media probes (red circles in Figure 3) are placed as containers running on Host 1 and Host 2 and interleaved with media processing units. The media processing units and probes are configured by the policy such that the video signals traverse the probe while going from one media processing unit

to the following one in the chain. Alternatively, the probe does not need to be directly in the chain but can be configured only to observe "tapped" (e.g., copied from a virtual switch) "mirrored" (e.g., copied from a physical switch port), or "captured" (e.g., collected from broadcast or multicast) media traffic.

In some cases, the media processing functions may also be configured to replicate the signal and send one signal replica to a virtual probe while the other replica is sent to the next processing function in the pipeline. This enables connecting a virtual probe to a processing function without requiring additional plumbing. In this particular case the communication between the processing function and the virtual probe can use different mechanisms such as shared memory, etc. The type of information a virtual probe provides may therefore depend on how the virtual probe is connected to the different probe points. For instance, information such as inter-packet gap cannot be provided when the communication between the virtual probe and the processing function is not network based.

The probes may also be configured to collect metrics of interests and to send them to a local and/or remote collector(s) (not shown in Figure 1). The metrics may include inter-frame time interval, frame timestamps, frame rate, number of blank frames, number of missing frames, etc. In addition, a virtual probe may be connected to several probe points to enable, e.g., measuring the phase difference between different streams, etc.

Second, two network switch based media probes (yellow stars in Figure 3) enable specific measurements in the switch ports traversed by the media signal on the specific media (IP) flows of interest. Figure 4 below illustrates a sample data center network topology. The solution descriptor can include configurations to be enforced on the ports highlighted in red. These configurations ensure that (1) traffic going from Host 1 and Host 2 towards Host 3 matching the IP tuple in use for the media signal under measurement is analyzed; and (2) analysis results are sent to a collector (not shown in figure). This mechanism can be seen as media-specific flow where matching packets are analyzed specifically for the media at hand. The metrics may include inter-packet delay, packet size, number of lost or dropped packets, and/or, if media specific functions are available, inter-frame delay, frame timestamps, frame rate, etc. Alternatively, the switch port can be configured only to mirror traffic of interest to an external analyzer.

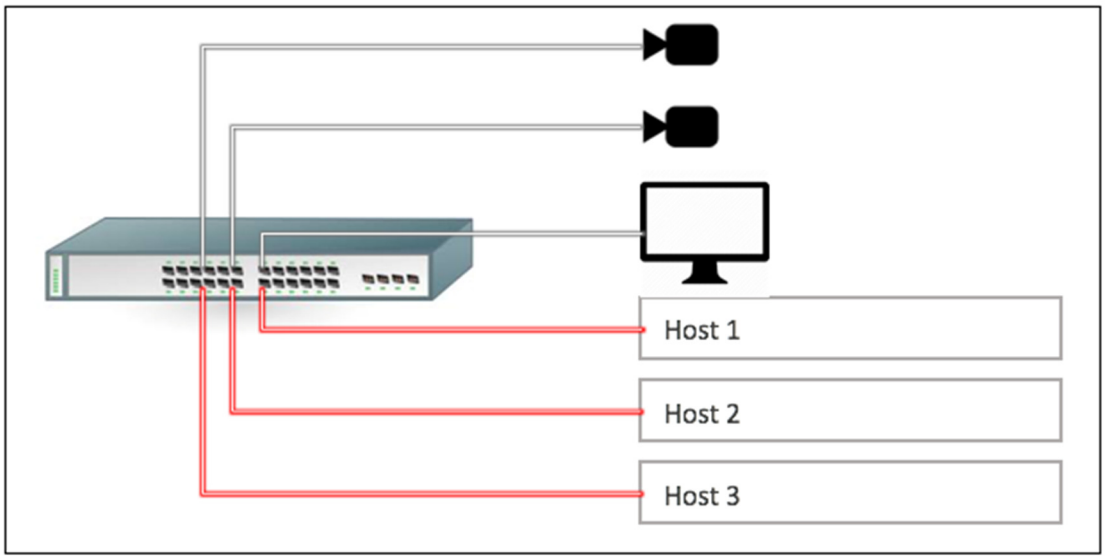

Figure 4: Sample media application - sample network topology

Third, one physical measurement device (orange box in Figure 3) is used to intercept and analyze traffic going between the bare metal processing unit and the media sink (depicted as a display). This consists of configuring the SDN controller, or the network devices in general, to funnel traffic matching the IP tuple of the media signal towards the physical measurement device. For example, Figure 5 below illustrates a network topology of a network that can be configured to divert media flow packets through the physical measurement device, as highlighted by the yellow thick line. Techniques that could be used include, for example, tunneling, "Virtual Local Area Network (VLAN) stitching," routing rules, etc. The physical measurement device is also configured to collect statistics of interest of the diverted flow and send results to a collector (not show in figure).

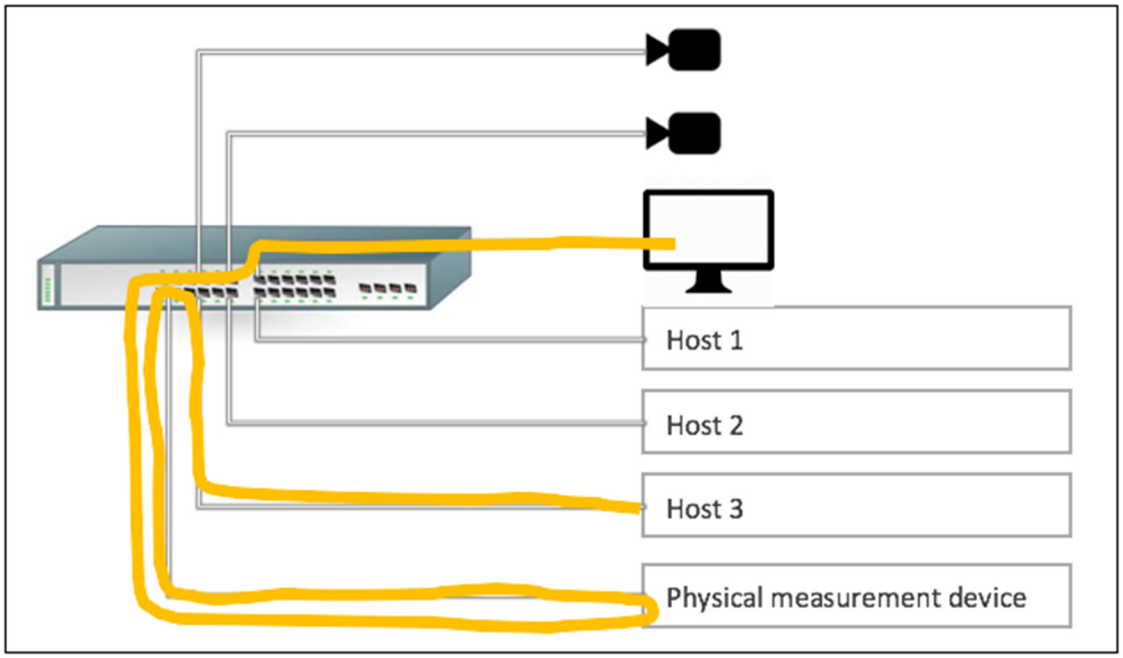

Figure 5: Sample media application - Funneling traffic

This example is just an instantiation of the policy, which thus can be applied to any media solution and can lead to a number of different resulting probing infrastructures. In particular, virtual probes can be implemented as containers, virtual machines (for different hypervisors), server-less functions (i.e., as code provided to a platform-as-a-service system), etc. The policy is aware of what is available and can thus decide accordingly.

Policy decisions are driven by the capabilities present in the datacenter (available physical measurement devices or switch features) and by policy configurations. For example, one configuration may request to use virtual probes wherever possible as opposed to physical ones, while another configuration may request to place probes at all pipeline stages or only in "important" points of the pipeline (for any given definition of importance). Yet another configuration may request to prefer a "funneling" method, while still another configuration may determine which particular metrics are activated, etc.

The techniques presented herein focus on media application, but some can be adopted in other domains as well, wherever traffic analytics might be beneficial. The value resides in providing a fine-grained, media-specific monitoring system for cloud-based media application, which may be more powerful than what currently exists in physical

media deployment. In addition, a media solution can be automatically produced through a policy driven orchestrator.

Summary

The techniques presented herein may involve applying media specific probing infrastructure and leveraging a policy driven orchestrator to automatically produce media-specific workflows. With respect to the policy driven orchestration of measurement probes, the policy may take into account many factors including user preferences, application type, and application topology to determine (1) the kind of probe and (2) where to place the probe. With respect to the orchestration of probes together with the application orchestration to support those probes, the system may be aware of the application and while deploying the probes, and also manipulate the application deployment to support such probes. In-line probes, for example, need the application component to be reconfigured to send the traffic to the probes.

The techniques presented herein also allow for the application of cloud generic measurement tools to the media use case. Network probes (measurement enabled switches) are commonly used to assess the health of a network infrastructure. A similar technique may be employed to assess the health of a media signal transported into a distributed cloud media application. For example, this may be used to augment switch capability to analyze media traffic, or to deploy measurement probes able to dissect media signal and assess its "health." There also exists support for hybrid physical and virtual probing infrastructure for media.

The techniques presented herein orchestrate (in a policy-driven fashion) the deployment of the measuring infrastructure itself: to decide which sensors to deploy, to decide where to deploy them, and then to actually deploy them.

Probes are orchestrated which are able to collect metrics in the application-specific domain (e.g., measuring the media signal). Hence, the deployment is driven by policies that are application aware. In particular, this takes into consideration the application types and application topology in order to make decisions regarding the kind of probes to deploy given the specific application component running and the capability of the underlying

cloud system and infrastructure and where to run those probes based on application topology.

Many probe types are orchestrated, and often such orchestration also impacts the way the application itself is deployed. For example, the orchestration system may need to reconfigure an application component (e.g., a camera or a video processing function) send data to the probe, reconfigure a switch to duplicate traffic, or instruct an SDN controller how to divert, route, copy, etc. monitored traffic where needed.

The techniques presented herein provide a deployment of fine-grained application-aware policies given the specific distributed application. This approach covers different kinds of probes (i.e., measuring devices) that need to be deployed and orchestrated in a cloud environment. The policy that decides what and where to deploy the probe is also application aware and thus not based only on network topology.

In summary, techniques are provided to measure and analyze the "health" of a cloud-based professional media application. In particular, presented are (1) a (media) probing infrastructure, based on a mix of virtual and physical measurement probes or collectors; and (2) a set of techniques that allow an orchestration system to control the probing infrastructure, to deploy and/or configure measurement probes, and/or to divert or mirror traffic to be "measured" through such probes. This approach not only uses virtual or physical probes, but also the automatic and policy-driven orchestration of such probing infrastructure and leverages IP probing and measurement mechanisms for the media space.