SCTE·ISBE CABLE-TEC EXPO 2017
DENVER, CO
OCTOBER 17-20

2017 Fall Technical Forum
SCTE·ISBE • NCTA • CABLELABS

Cloud-DVR Real-Time Splunk-Based Monitoring and Alerting System

A Technical Paper Prepared for SCTE/ISBE by

Shlomo Ovadia, Ph.D.
Director of CPE Engineering
Advanced Engineering, Charter Communications
14810 Grassland Drive, Englewood, CO 80112
720-279-2875
Shlomo.ovadia@charter.com

Jenson Thottian
Splunk/Dev-Ops Architect
Charter Communications
6380 S. Fiddlers Green Cir, Greenwood Village, CO 80111
720-721-2875
c-jenson.thottian@charter.com

Table of Contents

List of Figures

List of Tables

Introduction

Cable operators are responding to their subscribers' insatiable appetite for TV programming with new innovative video solutions such as Cloud-DVR (cDVR) and TV Everywhere [1, 2]. cDVR solutions move in-home recording and playout functions to the cloud, and thus enabling remote access of the recorded Linear TV content on a variety of platforms such as laptops, tablets, cellphone, and TVs. cDVR has many advantages such as cost, performance, operation, and business intelligence compared with in-home DVR. For example, no truck-rolls are needed to deploy and fix cDVR issues. The cDVR service offers a virtually unlimited number of tuners, enabling customer to record more than two shows at a time, and providing completely scalable and redundant storage capability. This feature allows customers to easily increase the amount of paid storage without any changes to their home network. Furthermore, cDVR based solutions allow cable operators to deploy faster, and to use more cost-effective CPE devices such as Charter's Worldbox 2.0 with traditional and cloud-based interfaces [3].

However, unlike the in-home DVR use case in which all the time-shifted recorded content is served locally, and thus has no impact on the cable network, cDVR has several infrastructure performance costs to consider. The cDVR use case requires network capacity on the cable access network, and storage capacity for the recorded content. The cDVR solution utilizes unicast video delivery for each subscriber, the total required network capacity is proportional to the number of concurrent cDVR subscribers viewing the time-shifted content. One of the potential cDVR obstacles is the copyright challenge by content providers. In the Private Copy deployment model, the cDVR permits each subscriber to record, store and playback a private and unique copy of the selected content (e.g., Private Copy). This means the cDVR storage capacity is linearly proportional to the number of subscribers. Another deployment option for cDVR is the hybrid storage model where the Private-Copies are maintained for 3 days after the record-time [4]. If the Private Copies are not viewed by then, they can be deleted ("de-duplicated") by the cDVR system and saved only as a Shared-Copies (same recorded content is shared among multiple subscribers). Thus, in this model, the amount of required storage for each subscriber can be significantly reduced compare with the Private Copy model.

Cloud-DVR System Architecture

Figure 1 shows Charter's cDVR system architecture. The main building blocks are the Video Storage and Processing Platform (VSPP) storage nodes, VSPP Manager's nodes, Scheduler servers, Geo-fencing servers, and KUMO servers and software. There are 56 VSPP storage nodes organized in four virtual entities called PODs. Each POD clusters 14 storage nodes using Commercial Off-The-Shelf (COTS) servers interconnected via a LAN topology using 1Gigabit Ethernet (GbE) interface for management, and 10GbE full-meshed inter-connection networks using high-speed switches. All the nodes in the POD contribute their physical resources in terms of storage capacity, CPU power, ingest and streaming throughputs. The VSPP software stack includes virtualized Software-Defined Storage (SDS), which enables high IO performance, high-availability storage solution with seamless fault-tolerance and self-healing operation. The VSPP SDS provides distributed Redundant Array Independent Disk (RAID) 5 storage used for content redundancy as protection against data loss due to disk failures. In addition, each storage node is connected to high-speed

Leaf switch using a 40GbE link, while each Leaf switch is connected to a high-speed Spine switch using a 100GbE link.

All the storage nodes run a Linux-OS and VSPP software stack, and are responsible for the actual data ingest and processing (e.g., transcoding, packaging, Ad-insertion), storage and streaming of content. The VSPP node joins the multicast stream for each linear feed to be recorded. Recording can only be performed for streams that have been requested by at least one user. All the Adaptive Bit Rate (ABR) video streams are delivered to using IP unicast with bi-directional TCP/IP connection between the ABR client device and VSPP storage node. All ABR video profiles for video and audio are stored up to 12 Mbps aggregated throughput. The cDVR peak storage capacity is 13.8PB. The VSPP Manager controls and orchestrates the entire VSPP storage nodes' activities and flows. The Scheduler is responsible for scheduling all the recordings based on the latest received EPG ingest information. KUMO is Charter's abstraction-layer software running on dedicated COTS servers, which receives information logs from client devices and manages and executes the different transactional REST APIs among the Scheduler, client devices, and the other back-office services such as IPVS, security system, NNS, etc.

Since Charter network spans a large portion of the U.S.A, it divides the U.S.A into disjointed geographical areas. Each subscriber has a unique HomeID, and an array of geo-fencing feed tags. A logical channel has different market-level feeds, which are combinations of a unique feed identifier (UUID) and channel name, broadcasted on different geographical regions defined as a list of geographical identifiers (geoID). Not all the logical channels are available over the entire US geo areas. When recording a channel having multiple market-level feeds, the VSPP system must record the market-level feed that applies to the location of the subscriber, which is provided by the Geo-Fence service. There are two Geo-Fencing servers configured in active/passive architecture.

In addition, the VSPP platform includes a Diagnostic server, which is a centralized monitoring and diagnostic back-end server. It continuously collects and aggregates all the VSPP system health and service-level metrics, and sends SNMP traps and Splunk data logs to the cDVR Splunk real-time monitoring tool.

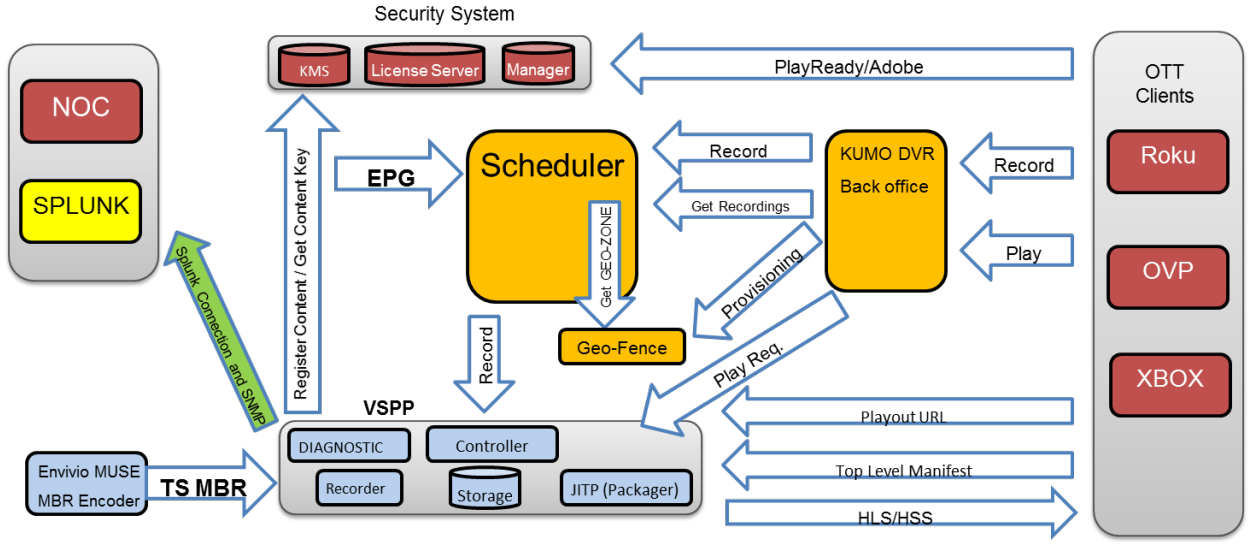

Figure 1 - Cloud-DVR System Architecture Showing the key Building Blocks and Interface to Splunk Environment

Real-Time Splunk-Based Monitoring & Alerting System

1. Monitoring Tool System Architecture

A Splunk-based real-time monitoring and diagnostics tool for cDVR service in production environment was developed. Figure 2 shows, for example, a high-level block diagram of the CDVR main dashboard with all the key components, including the Scheduler and Geo-Fence servers and apps, KUMO servers and apps, the VSPP nodes and apps, VSPP Manager, and Client devices. Two Splunk heavy-forwarders are forwarding VSPP metrics, application logs, Session Data Reports (SDRs), and SNMP traps to the Splunk indexer. In addition, KUMO servers' health metrics and application logs as well as KUMO API analytics, and application logs from Client devices such as OVP and Roku are being forwarded to the Splunk Indexers. All the background jobs are running Splunk ad-hoc searches, and sending all the collected metrics with in-house developed app to Charter's Splunk monitoring tool dashboard. Notice that the number of key components and metrics associated with each component is scalable, depending on the cDVR system complexity, to enable adequate video operation support. Furthermore, the main dashboard includes the following information:

- List of critical issues received in the last one hour
- List of SNMP traps received from the Diagnostics server in the last 4 hours according to their severity level (e.g., the highest-severity traps show-up first)
- Key cDVR plots such as total and used storage status
- cDVR metric definition and threshold levels
- Link to general cDVR analytics with a selectable-time period such as:
    - Number of daily/weekly users as reported by KUMO
    - Number of daily/weekly single and series recordings
    - Number of daily/weekly single and series playback recording
    - Number of daily/weekly recording failures
    - Number of daily/weekly deleted recordings
    - Top 10 watched channels

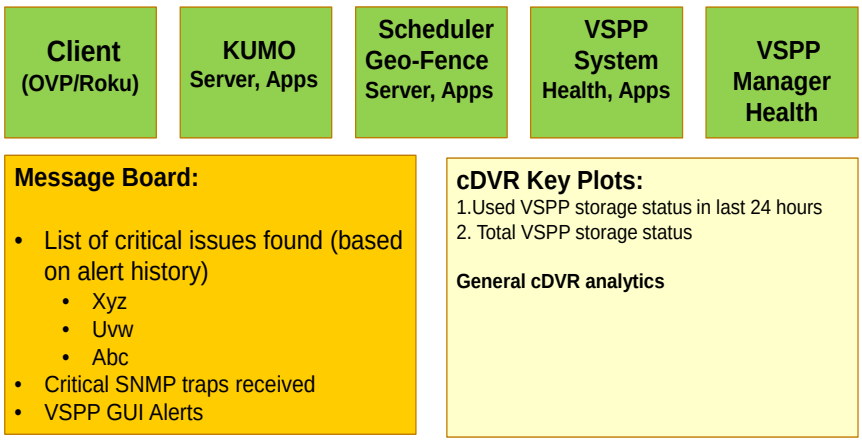

Figure 2: High-Level Block Diagram of Splunk-Based Monitoring Tool Dashboard

Figure 3 shows the main cDVR system health dashboard with its main subsystems, level 1 metrics for each of the subsystems, and their operational status. The listed subsystems and level 1 metrics are shown as an example, and other subsystems such as Arista switches and/or level 1 metrics are planned to be added. KUMO, which is a software abstract layer, receives information logs from client devices and manages and executes the different transactional REST APIs among the Scheduler, Geo-Fencing servers, client devices, and the other back-office services such as IPVS, NNS, etc. The VSPP Manager controls and orchestrates the entire VSPP nodes' activities and flows.

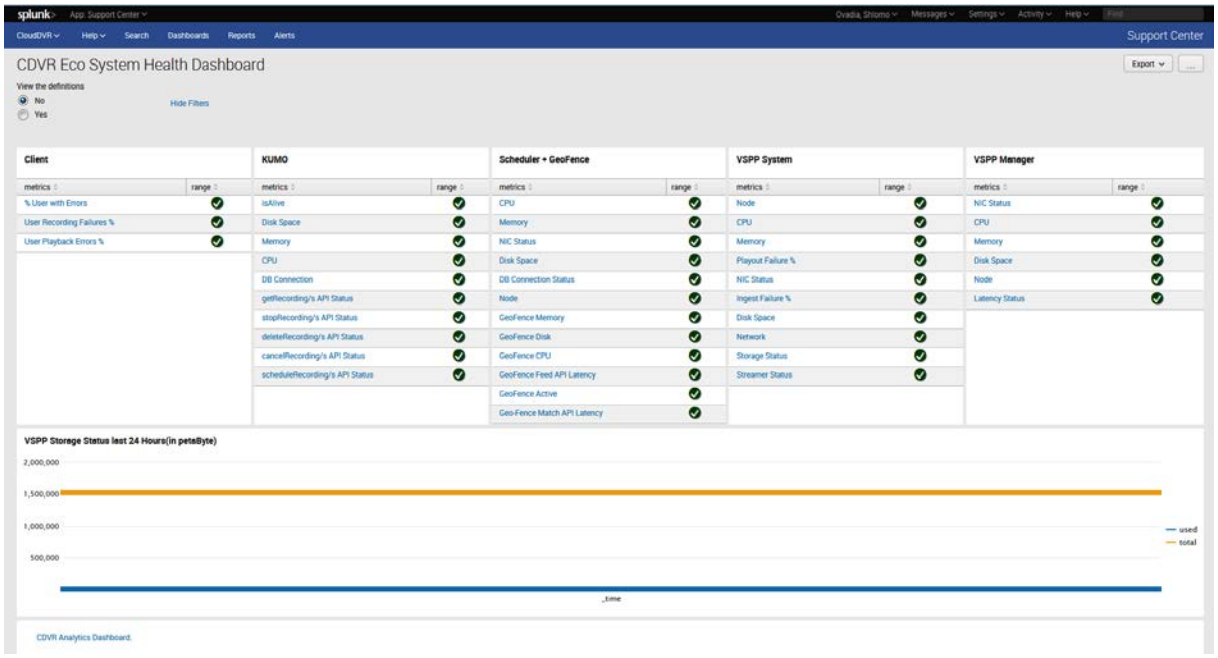

Figure 3: cDVR System Health Dashboard Showing All the Main Components, Metrics, and KPIs

Table 1 - provides a detailed description of the main cDVR dashboard components.

Table 1: cDVR Dashboard Component and Metric Description

cDVR Component	Description
VSPP Manager	Provide real-time level 1 metrics on: • Health of the VSPP Manager – CPU and memory usage averaged over 15 minute period, most recent available disk capacity, and how long the server has been up and running. • App analytics such as round-trip latency between the VSPP Manager and the selected storage node for the recorded sessions

	• Maintenance – how many storage nodes are going through maintenance operation such as software updates
VSPP System	Provide real-time level 1 metrics on: • Health of each of the 56 storage servers, including CPU and memory usage, most recent available storage disk space, and how long each server has been up and running • Network status – total number of errors while either receiving or transmitting packets • Storage Throughputs – disk write and read throughputs on each VSPP storage node • App analytics – relative performance of various app running on the VSPP system, including live ABR ingest and playout failures (%) in 60s period and round-trip latency of closed recording for each ABR session • Streamer status - provide activity status (active/not active) of the processes running on each of the storage nodes
Scheduler Server	Provide real-time level 1 metrics on the health of the Scheduler servers, including CPU and memory usage averaged over 15 minute period, most recent disk space utilization, and how long the node has been up and running
Geo-Fence Server	Provide real-time metrics on the health of the Geo-Fencing servers, including CPU, memory usage, disk space, NIC status, node status, and the following API analytics: • Geo-Fence Match API latency – round-trip latency between the Geo-fence server and its REDIS database • Geo-Fence Active – identify which Geo-fencing server is currently active • Geo-Fence Feed API latency – round-trip latency between Geo-fence server the KUMO server
KUMO Server	Provide real-time level 1 metrics on: • Health of the KUMO servers - CPU and memory usage averaged over the last 15 minute period, available disk capacity, and how long the server has been up and running. • DB and VSPP Connectivity – are the KUMO servers connected to the VSPP system and its own database • isAlive – Check if the KUMO app (Java) is alive and running • KUMO API analytics – obtain analytics such as unsuccessful transaction (%), error rate, transaction duration for various APIs: o getRecording(s) o stopRecording(s) o cancelRecording(s) o scheduleRecording(s) o deleteRecording(s)
Client	Provide real-time level 1 metrics on user recording and playback errors (%) and users with errors (%) for various client devices such as OVP and Roku.

2. Dashboard Features and Health Metrics & KPIs

The cDVR system dashboard health metrics and Key Performance Indicators (KPIs) are organized in a hierarchal fashion to facilitate the consumption of the vast amount of available information. When each of the level 1 health metrics and/or KPIs status changes to either a warning or critical condition (orange or red symbol), an operational engineer is able to obtain further detailed information as shown in Table 2.

Table 2: Hierarchal Organization of cDVR System Health Dashboard Metrics

Metric Level	Description
1	Metric status information based on pre-defined threshold levels for each CDVR subsystem
2	Metric status information such as the hostname of the server, its health status based on predefined threshold levels, name of transactional APIs and their status based on error count
3	Time-based behavior of the selected Level 2 metric
4	Event-based result showing the detailed information of the selected level 3 metric at a specific time.

Figure 4 shows an example of VSPP system health dashboard with level 2 metrics such as CPU and memory usage average over 15 minutes period, and disk space utilization on each of the displayed storage nodes. For each level 2 metric, the dashboard shows the hostname of the node, which POD it belongs to, and the corresponding metric health status based on pre-defined threshold levels. In addition, the storage disk utilization (in %) for each POD is displayed. If a storage node undergoes a maintenance operation or has failed, it will be indicated in the level 2 Node Status metric. The number of failed storage nodes is simply equal to 56 – (number of active nodes) – (number of nodes undergo maintenance).

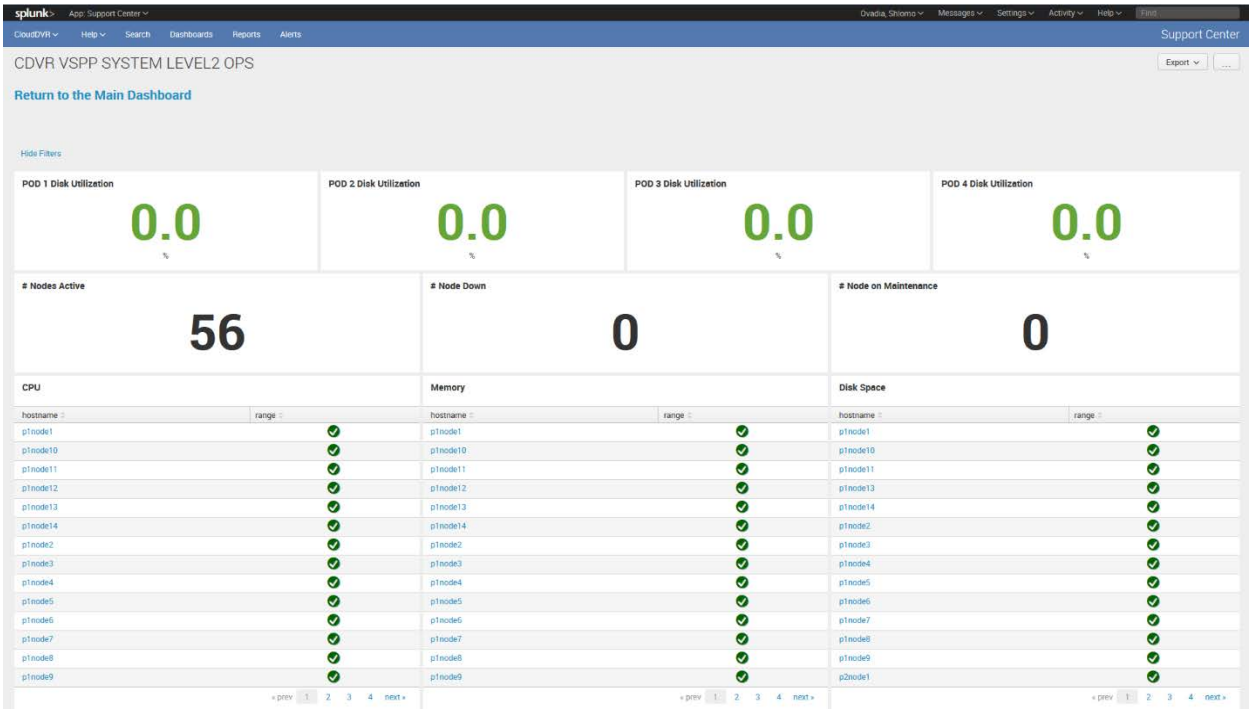

Figure 4: VSPP System Health showing POD Disk Utilization, and Level 2 Metrics such as Disk Space, CPU and Memory Usage for each Storage Node

Level 3 metrics provide the operational engineer with the level 2 metrics behavior versus time. Each level 2 metric can be filtered and displayed as follows:

- Time-base filtering: operational engineer can select the specific-time duration of interest to view the selected metric behavior
- Hostname filtering: operational engineer can select the specific hostnames to be displayed on the dashboard in the selected time frame.

Figure 5 shows, for example, the CPU and memory usage (%) of p1node12, p2node12, and p3node12 in the last 24 hours. This type of dashboard display is particularly useful for a time-dependent metric comparison among different storage nodes. If there are any storage nodes in maintenance, then the top part of the level 3 metric dashboard will show the hostname of each storage node in maintenance. Splunk query information is available by clicking on the observed trace for a specific metric on the selected hostname and time duration. Each of the dashboard for level 1, 2, and 3 metrics is updated every 15 minutes. Since there are so many concurrent background jobs, it takes two 15 minutes periods to update every dashboard metric. The 15 minute period was selected as a design trade-off between the number of concurrent background jobs for a dashboard update and the rate in which changes occur in the cDVR system.

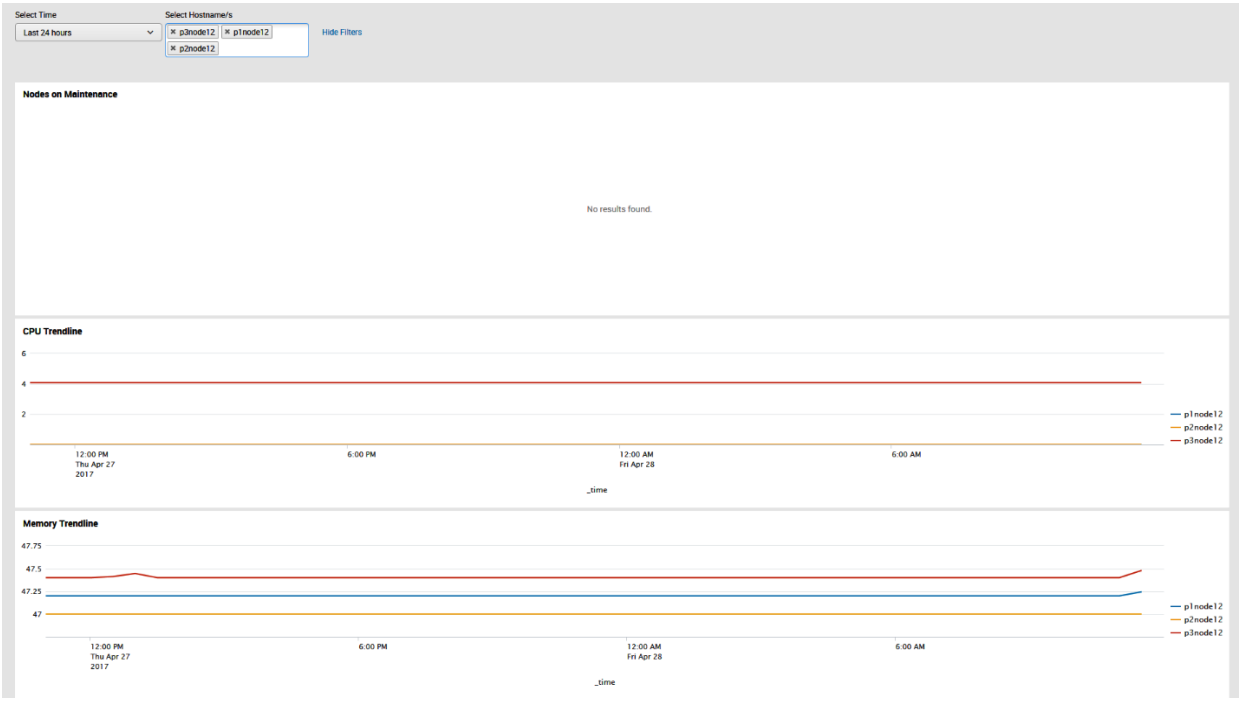

Figure 5: CPU and Memory Usage (%) on p1node12, p2node12, and p3node 12 in the last 24 Hours

3. Server API Monitoring and Alerting

Various cDVR server transactional APIs such as from KUMO or Geo-Fencing servers are being monitored on the dashboard. Since the application logs are ingested within the Splunk environment, it allows the user to obtain valuable analytics about the monitored APIs. This includes detailed information about various transactions occurring in the VSPP system such as:

- Transaction count and duration in the specific time period
- Average transaction duration
- Transaction error rate (%) in the selected time period
- Percentage of unsuccessful transactions in the selected time period

Figure 6 shows, for example, the getRecordings API transaction status in the last 30 days. The getRecordings API transaction represents all series recordings requests sent by KUMO to the Scheduler. Table 3 summarizes the various reported metrics for getRecordings API. Notice that the getRecordings API transaction duration has large time variations, and 4.267 % of all the transactions were unsuccessful since these transactions timed out. The API transactions were timed out since their duration exceeded the one second threshold level. This programmable threshold level allows the cable operator to make sure that the cDVR system performance is within its design boundaries.

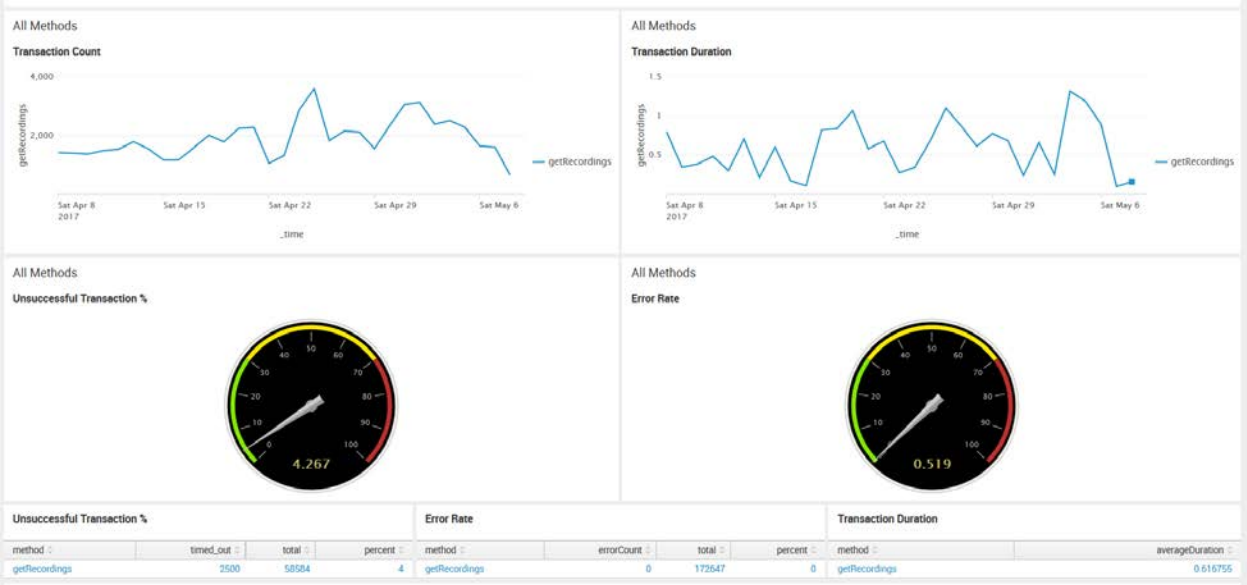

Figure 6: KUMO getRecordings API Status in the Last 30 Days

Table 3: Reported getRecordings API Transaction Metrics

Reported Metric	Reported Value
Peak API Transaction Count	3593
Average Transaction Duration	616.755ms
Transaction Error Rate	0.519%
Unsuccessful Transactions	4.267%

A key feature of the Splunk-based monitoring dashboard is the ability to change the threshold levels for each metric as needed. Programmable threshold levels are defined for each metric to indicate a healthy node condition, a warning or a critical condition. When the status of each metric changes to either a warning or critical condition, the operation engineer is able to obtain more detailed information via level 2 or level 3 metrics. This feature enables threshold-level adaptation of each metric based on cDVR system behavior in production environment.

Another key feature of the Splunk-based dashboard is the ability to send alert notifications and customized detailed reports via e-mail to a group of users when each of the KPI or level 1 health metrics crosses a critical threshold level. Specifically, e-mail notifications are sent when the status of any of the level 1 metrics or KPIs is changed from green to red. Furthermore, e-mail notifications will continuously be sent every hour until the red metric status is back to a healthy condition. With the addition of diagnostics capabilities, the operational engineer is able to take the necessary steps to resolve the observed issues.

4. SNMP Traps and Messages

The Splunk-based monitoring tool receives SNMPv3 traps and GUI messages from various cDVR components that are not shown in Figure 2. Table 4 lists other SNMP traps and messages that are received by the cDVR system dashboard. Both the health status of the high-speed network switches and the low-level server's hardware alerts information is available on level 2 of the VSPP system dashboard.

Table 4: Other SNMP Traps and Messages Received by cDVR System Dashboard

cDVR Component Name	Description	Provided Information
Network Switches	High-speed Spine and Leaf network switches connecting the VSPP storage and management servers	• Deep insight and visualization of Spine and Leaf switches' health metrics • SNMP traps to identify any failures and performance degradation.
Server Hardware Alerts	Low-level storage and management servers' hardware alerts	Critical and major server's hardware fault events such as power supply, memory, storage disk, fan or smart battery array failures as defined via the vendor's SNMP MIBs.
VSPP GUI Messages	VSPP system syslog messages	VSPP syslog messages about various VSPP internal configuration setting changes listed according to their severity level and received time.

5. Diagnostics Features

A real-time list of new critical or warning issues are generated by cDVR system alerts and received SNMP traps based on pre-defined threshold levels for each metric. For rapid and scalable diagnostics of the observed issues, the monitoring tool needs to guide the operational engineer with an action plan in order to resolve the observed critical issues. Two primary diagnostics capabilities are proposed as follows:

A. Metric Diagnostics and Escalation Path

When the status of any level 1 metrics changes to either warning or critical, a new diagnostics screen becomes available. This new diagnostics screen is in addition to the level 2 metrics screen, and occurs when the operational engineer clicks on the selected level 1 metrics. The diagnostics screen provides the following capabilities:

- Suggested list of steps based on all the received system alerts and SNMP traps for the operational engineer to check in order to address the observed warning or critical level 1 metrics.
- Escalation path with contact info (e-mail address and phone #) of advanced engineers for further debugging of the observed issues.

B. Knowledge Based Diagnostics

Based on the video operation experience of the cDVR system, it is expected that some of the observed failures or critical issues may not be immediately solved by the operational engineers. A potential solution may be found based on the generated application error codes and their severity. When an observed failure, warning or critical issue is resolved, a detailed report may be generated with the suggested best engineering practices how to address such a future failure or critical issue. Depending on the Root-Cause Analysis (RCA) of the observed failures and the number of impacted customers, the threshold levels of some level 1 metrics may be changed, and/or new metrics maybe added. For example, after the Scheduler software was recently upgraded, it was found that the Scheduler had a stale EPG data, resulting in many failed recordings. Based on the vendor RCA, a new EPG ingest status metric that checks if the Scheduler has the latest EPG data is being added to the Scheduler dashboard.

Such knowledge-based reports are converted into a list of steps for the operational engineers to check before escalating to Tier 2 or 3 support engineers. Thus, the integration of this type of knowledge-based diagnostics capabilities into the monitoring tool can significantly reduce the time it takes the operational engineers to resolve new failures, warning or critical issues.

6. Self-Learning Monitoring Features

Self-learning monitoring capabilities are essential in order to continuously improve the monitoring tool. Three different types of self-learning capabilities are identified as follows:

1. **Tool health Check:**
   The tool has a shell test script that periodically runs and checks if any background jobs are not running and reporting the assigned metrics. If the test script finds one or more such background jobs that didn't run, the test script performs the following steps:
   a. Check the status of all the background jobs that were run in a last period of time (i.e., 15 min.) according to their priority, and
   b. Rerun all the jobs whose test result was inactive (e.g., these jobs didn't run)
      i. Check if this issue previously reoccurred with specific background jobs. If yes, the test script assigns these jobs a highest priority when executed. Test script monitors if no issue reoccurred, for example, in a four-week period, then the test script reduces the specific background job priority.

   This feature allows the operational engineer to monitor that the health of the monitoring tool.

2. **Self-Optimization of Metric's Threshold Levels:**
   By accessing historical Splunk logs and observing time-dependent behavior patterns of various metrics (level 3) in different dashboard components, the monitoring tool automatically reprogram the pre-defined initial threshold levels of the metrics, and provides an updated report on the updated threshold levels for each metrics. The tool provides a gradual incremental change in the pre-defined threshold to avoid false reporting of system health. In addition, this allows the operational engineer to reject some of the tool's reprogram threshold levels, and enter new threshold levels for the specific metrics and/or KPIs.

3. **Operation Intelligence:**

The VSPP Manager, which control and orchestrate all the activities in the cDVR system, monitors the health of each storage node as well as the apps that are running on the node. If, for example, a storage node starts to exhibit hardware failures or an abnormal behavior such as:

- One or more storage disk failures as received by an SNMP trap
- High-temperature inside a storage node as received by the iLO alerts
- Very high CPU or memory utilization (> 90%)

In this case, the monitoring tool sends a REST API request command to the VSPP Manager to take this storage node in a maintenance mode. The VSPP Manager put the specified storage node in a maintenance mode, and sends REST API acknowledgement to the tool. This allows the operational engineer to further debug the issue and take the appropriate action such as:

- Replace the failed storage disk
- Replace other parts within the server
- Reboot the node to check if the node is reporting healthy behavior, or
- Perform a scheduled software update/upgrade

After the hardware failure is fixed or the software upgrade is completed, the monitoring tool performs the following tasks:

- Checks the status of the following metrics and alerts:
  - All the reported level 2 metrics are healthy from this storage node
  - There are no SNMP alerts from the Diagnostics server
- If yes, it sends a REST API request to the VSPP Manager to take off the storage node from maintenance mode
- Generate a detailed report of the incident

The VSPP Manager takes the node off the maintenance mode, and completes the request by sending a REST API acknowledgement to the monitoring tool. By logging these cases, the monitoring tool can take similar actions if hardware or software failures re-occurred in other nodes.

Comparison with Other Monitoring Tools

There are other enterprise-level monitoring tools that can be used to monitor the cDVR system. Table 5 shows a high-level comparison between Splunk Enterprise [5], Graphite/Grafana [6], Nagios XI [7], and ELK [8] monitoring tools. The Splunk Enterprise is a flexible and scalable platform that makes it simple to collect and analyze vast amount of machine data, and act upon the received system alerts and SNMP traps. However, this is a proprietary monitoring solution that requires customer subscription. Fortunately, there are many vendors that already developed Splunk-based telemetry applications for their systems, which simplifies the integration of these applications into a Splunk-based dashboard.

Graphite is an open-source app for collecting, analyzing, and providing real-time monitoring of server health and application metrics for enterprise platforms. In addition, it offers a user-friendly graphical presentation of the data via the Grafana dashboard. However, the Grafana dashboard doesn't receive SNMP traps, and doesn't send e-mail notification or has an ability to set-up threshold levels for various metrics, which limits its usefulness in a large-scale video operation.

Table 5: Comparison between Splunk, Graphite, and Nagios XI Monitoring Tools

Monitoring Tool	Pros	Cons
SPLUNK>	• Enterprise-class high availability • Scalability • Customized dashboard • Interactive graphs • Server and App metrics • Ability to receive SNMP traps • Ability to send e-mail notifications • Generate customized technical reports • Data logs retention & reporting • Provide analytics for Nagios XI • Specialized modules are available for security, IT services, and user behavior • Easier to use logs for troubleshooting	• Proprietary • Required customer subscription • Potentially expensive
Graphite/ Grafana	• Open-source tool • Interactive graphs • Scalability • Low-cost subscription	• Doesn't receive SNMP traps • No ability to set threshold-levels to dashboard metrics • No alerts or e-mail notifications • No or limited technical support • Limited data logs retention • No support for string metric values
Nagios XI	• Open-source tool • Customized dashboard • Scalability • Server and App metrics • Ability to send e-mail notifications • Generate technical report • Data logs retention and reporting	• No ability to set threshold-levels to dashboard metrics • No ability to receive SNMP traps • GUI lacks user-friendliness • Requires subscription for Enterprise-level tool (potentially expensive)
ELK	• Open-source tool • Scalability • Customized dashboard • Server and App metrics • Data logs retention and reporting • Alerts or e-mail notifications (with X-Pack) • Generate technical report	• Missing user management features (in basic ELK) • No SNMP traps (w/o using external modules) • ELK cluster deployment requires more time & resources than Splunk • Data onboarding is harder than Splunk • Feature-poor UI compared with Splunk • Only accept JSON-formatted data • No specialized modules are available for security, IT services, etc.

Nagios XI provides monitoring of all mission-critical infrastructure components including applications, services, operating systems, network protocols, systems metrics, and network infrastructure. Hundreds of third-party add-ons provide for monitoring of virtually all in-house applications, services, and systems. Although Nagios is the best known free monitoring tool, its open-source version is limited in terms of dashboard features and capabilities, there is a steep and costly learning curve, and the GUI lacks user-friendliness. This is one of main reasons that the Nagios XI tool is losing its appeal among corporate customers.

ELK is an open-source monitoring tool that is gaining popularity among cooperate users similar to Splunk. It uses Elasticsearch for ingesting data logs, and Kibana as the visual UI for displaying customized dashboards. Enterprise users can purchase X-Pack, which is an Elastic Stack extension that bundles security, alerting, monitoring, reporting, and graph capabilities. X-Pack components are designed to work together seamlessly, allowing the user to enable or disable the desired features as needed. However, a larger ELK development effort for a customized monitoring solution than in Splunk may be needed, depending on the growth rate and complexity of deployment use cases. Another main difference is the way the data is parsed. ELK requires you to identify the data fields before it's shipped to Elasticsearch, while with Splunk, you can do that after the data is already in the system. This makes data onboarding easier by separating shipping and data classification/field labeling.

Summary

In this paper, the cDVR system architecture with its key hardware and software components was reviewed first. Then, the real-time Splunk-based cDVR monitoring system architecture, main features, health metrics, and KPIs were explained. This includes four-level hierarchal organization of server health metrics from VSPP storage nodes, VSPP Manager, Scheduler, Geo-Fencing and KUMO servers as well as all the transactional APIs' analytics from KUMO, Geo-Fencing, and VSPP Manager. The cDVR system dashboard includes the status of each health metric based on pre-defined programmable threshold levels for healthy (green), warning (amber), and critical (red) condition. When the status of each metric is changed from green to red, e-mail notifications and customized health reports are sent to the operational engineers to take the necessary actions to resolve the issues. A complete end-to-end system health report is achieved via the received SNMP traps by the cDVR dashboard from the high-speed network switches, low-level storage and management servers' hardware, and VSPP system syslog messages listed according to their severity level.

Another novel feature of the cDVR monitoring and alerting system is the self-learning capabilities such as the monitoring its own health. It periodically runs and checks if one or more background jobs are not running and reporting the assigned metrics. Furthermore, the monitoring system provides self-optimization of the threshold-levels for each of the monitored metrics and KPIs based on historically system behavior and observed failures. Another key aspect of the self-learning capabilities is the operation intelligence to identify low-level hardware failures such as a disk, fan, memory or smart array battery failures in a VSPP storage node, put the node in maintenance for further user diagnostics, and take the node off maintenance mode after either hardware repairs or software update tasks are completed.

Although the monitoring and alerting system cost is an important factor when comparing Splunk-based monitoring tool with other enterprise-class monitoring tools, the available features set, data logs ingest, scalability, user-friendly graphical interface, and time-to-market to develop a real-time operation-ready monitoring tool are also important considerations. When comparing the pros and cons of all these considerations (e.g., Table 5), the Splunk-based monitoring tool appears to be the most suitable for our

cDVR system. Furthermore, the Splunk Enterprise environment is conducive for integration with other Charter's back-office applications such as IPVS, NNS, etc., resulting in faster video operation readiness.

Abbreviations

Table 6: Table of Abbreviations

Abbreviation	Stand For
ABR	Adaptive Bit Rate
API	Application Programming Interface
cDVR	Cloud Digital Video Recorder
CPE	Customer Premise Equipment
COTS	Commercial-Of-The-Shelf
CPU	Central Processing Unit
EPG	Electronic Program Guide
GbE	Gigabit Ethernet
GUI	Graphical User Interface
iLO	Integrated Lights Out
IPVS	IP Video Systems
KPI	Key Performance Indicator
MIB	Management Information Base
NNS	National Navigation Services
NOC	National Operation Center
OVP	Online Video Platform
RAID	Redundant Array of Independent Disks
RCA	Root Cause Analysis
REST	Representational State Transfer
RS-DVR	Remote Storage DVR
SDR	Session Data Report
SDS	Software-Defined Storage
SNMP	Simple Network Management Protocol
VSPP	Video Storage and Processing Platform

Bibliography & References

[1] Carol Ansley and John Ulm, "The Dawn of Cloud-Based DVR Services", SCTE Cable-Tec Expo, October (2013).

[2] John Horrobin and Yoav Schreiber, "Unicast or Multicast for IP Video? Yes!", SCTE Cable-Tec Expo, October (2014).

[3] http://www.multichannel.com/news/content/charter-taps-arris-key-development-partner-worldbox-20/408379

[4] I. Tomer, Hybrid Solution for Cloud DVR: Meeting the Needs of Converging Legacy and OTT Platform, BEC proceedings (2016).

[5] https://www.splunk.com/en_us/products/splunk-enterprise.html

[6] Overview of Graphite tool can be found at https://graphiteapp.org/#overview

[7] https://www.nagios.com/products/nagios-xi/

[8] https://www.elastic.co/webinars/introduction-elk-stack

