

CABLE-TEC EXPO® 2017

SCTE • ISBE

THE NEXT BIG...

DEAL
CONNECTION
INNOVATION
TECHNOLOGY
LEADER
NETWORK



DENVER, CO
OCTOBER 17-20



VIRTUALIZING VIDEO AND CLOUD



SCTE · ISBE

Cloud-DVR Real-Time Splunk-Based Monitoring & Alert System

Shlomo Ovadia, Ph.D.
Director, CPE Engineering
Charter Communications

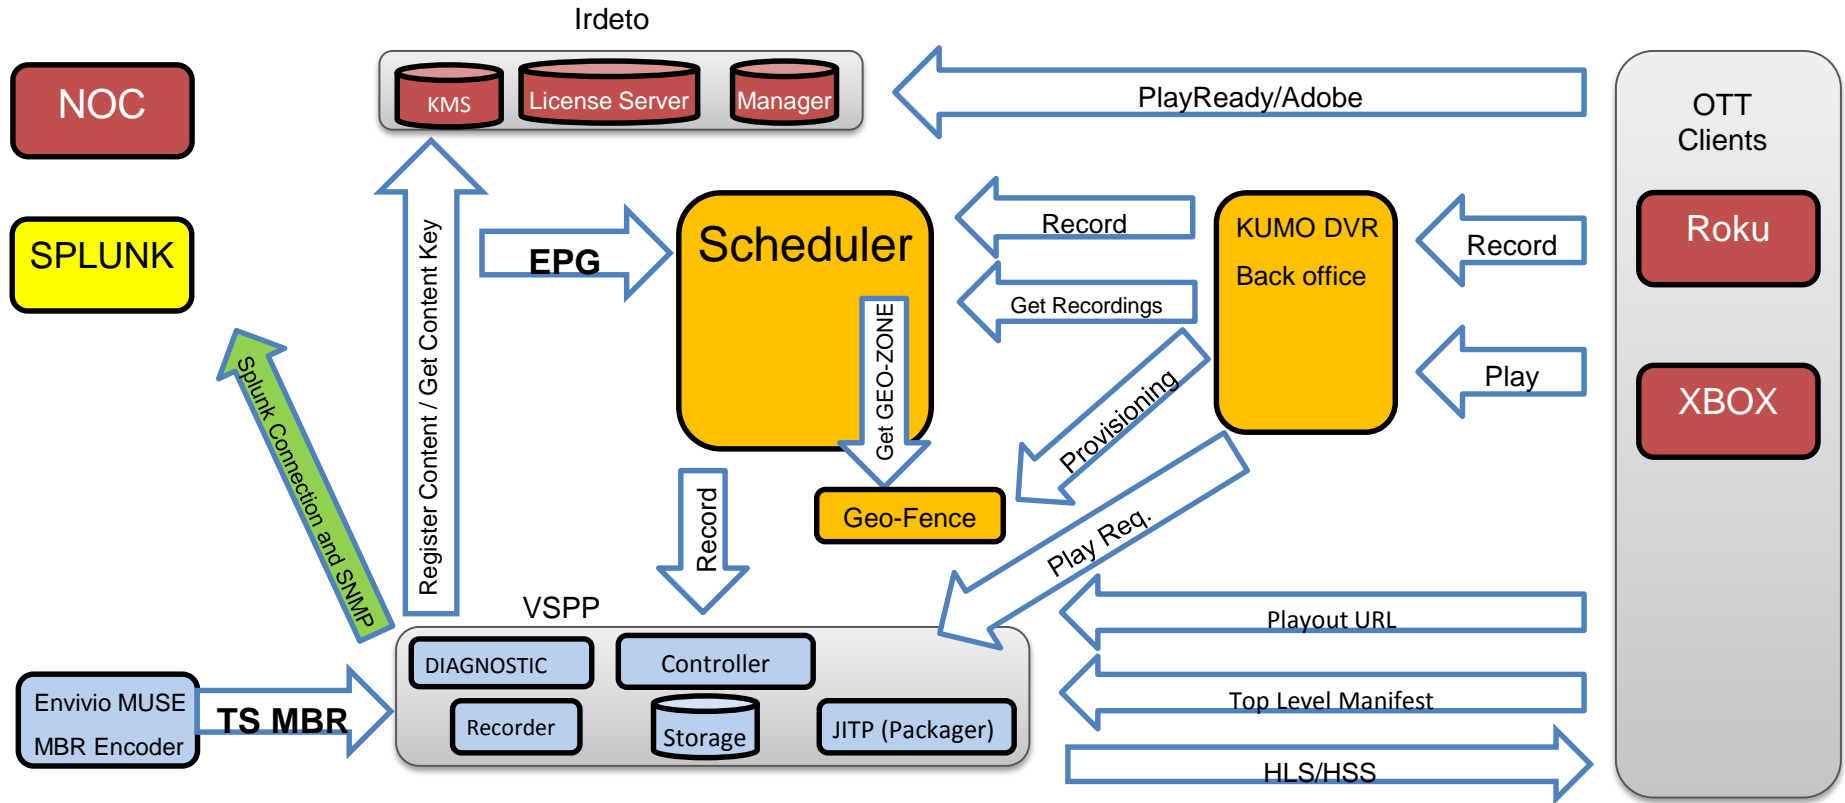


DENVER, CO
OCTOBER 17-20

OUTLINE

- Introduction
- Cloud-DVR (cDVR) System Architecture
- Splunk-Based Monitoring & Alerting System
 - Monitoring System Architecture
 - Dashboard Features & Health Metrics
 - Server & API Monitoring & Alerting
 - SNMP Traps and Messages
 - Diagnostics Features
 - Self-Learning Monitoring Features
- Comparison with Other Monitoring tools
- Summary

cDVR SYSTEM ARCHITECTURE



VSPN STORAGE NODES & MANAGER



- 56 VSPN storage nodes organized in four virtual entities called PODs.
- Each POD clusters 14 storage nodes using commercial off-the-shelf servers interconnected via a LAN topology using 1G 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.
- VSPN software stack includes virtualized Software-Defined Storage (SDS), which enable high IO performance, high-availability storage solution with seamless fault-tolerance and self-healing operation.
- VSPN SDS provides distributed Redundant Array Independent Disk (RAID) 5 storage used for content redundancy as protection against data loss due to disk failures.
- 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.

VSPP SYSTEM AND KUMO



- All the storage nodes run 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.
- 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 servers run the KUMO abstraction-layer software, 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, etc.

GEO-FENCING SERVERS



- 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.

MONITORING SYSTEM ARCHITECTURE



Message Board:

- List of critical issues found (based on alert history)
 - Xyz
 - Uvw
 - Abc
- Critical SNMP traps received
- VSPP GUI Alerts

cDVR Key Plots:

1. Used VSPP storage status in last 24 hours
2. Total VSPP storage status

General cDVR analytics

VSPP MANAGER and SYSTEM COMPONENTS



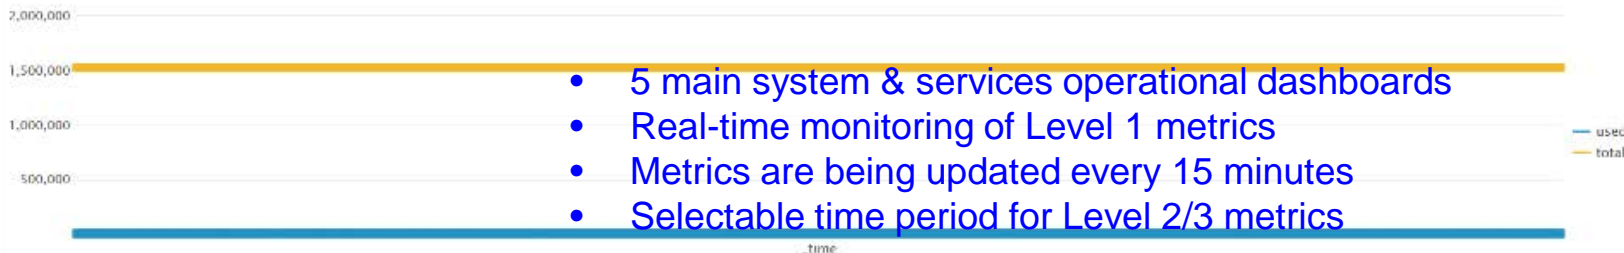
| cDVR Component | Description |
|----------------|---|
| VSPP Manager | <p>Provide real-time metrics on:</p> <ul style="list-style-type: none">• Health of the VSPP Manager – CPU and memory usage averaged over 15 minute period, most recent available disk capacity, and how long the server has been up 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 | <p>Provide real-time metrics on:</p> <ul style="list-style-type: none">• 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 |

| cDVR Component | DESCRIPTION |
|-------------------------|--|
| Geo-Fence Server | <p>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:</p> <ul style="list-style-type: none"> • 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 | <p>Provide real-time level 1 metrics on:</p> <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> ○ getRecording(s) ○ stopRecording(s) ○ cancelRecording(s) ○ scheduleRecording(s) ○ deleteRecording(s) |
| Client | <p>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</p> |

cDVR SYSTEM DASHBOARD

| Client | | KUMO | | Scheduler + GeoFence | | VSPP System | | VSPP Manager | |
|---------------------------|---------|--------------------------------|---------|-----------------------------|---------|------------------|---------|----------------|---------|
| metrics ▾ | range ▾ | metrics ▾ | range ▾ | metrics ▾ | range ▾ | metrics ▾ | range ▾ | metrics ▾ | range ▾ |
| % User with Errors | ✓ | isAlive | ✓ | CPU | ✓ | Memory | ✓ | CPU | ✓ |
| User Recording Failures % | ✓ | CPU | ✓ | NIC Status | ✓ | CPU | ✓ | Memory | ✓ |
| User Playback Errors % | ✓ | Disk Space | ✓ | Memory | ✓ | Ingest Failure % | ✓ | Disk Space | ✓ |
| | | Memory | ✓ | Disk Space | ✓ | Disk Space | ✓ | NIC Status | ✓ |
| | | DB Connection | ✓ | DB Connection Status | ✓ | NIC Status | ✓ | Latency Status | ✓ |
| | | getRecording/s API Status | ✓ | Node | ✓ | Node | ✓ | Node | ✓ |
| | | stopRecording/s API Status | ✓ | GeoFence CPU | ✓ | Payout Failure % | ✓ | | |
| | | deleteRecording/s API Status | ✓ | GeoFence Memory | ✓ | Network | ✓ | | |
| | | cancelRecording/s API Status | ✓ | GeoFence Disk | ✓ | Storage Status | ✓ | | |
| | | scheduleRecording/s API Status | ✓ | GeoFence Feed API Latency | ✓ | Streamer Status | ✓ | | |
| | | | | Geo-Fence Match API Latency | ✓ | | | | |
| | | | | GeoFence Active | ✓ | | | | |

VSPP Storage Status last 24 Hours(in petaByte)

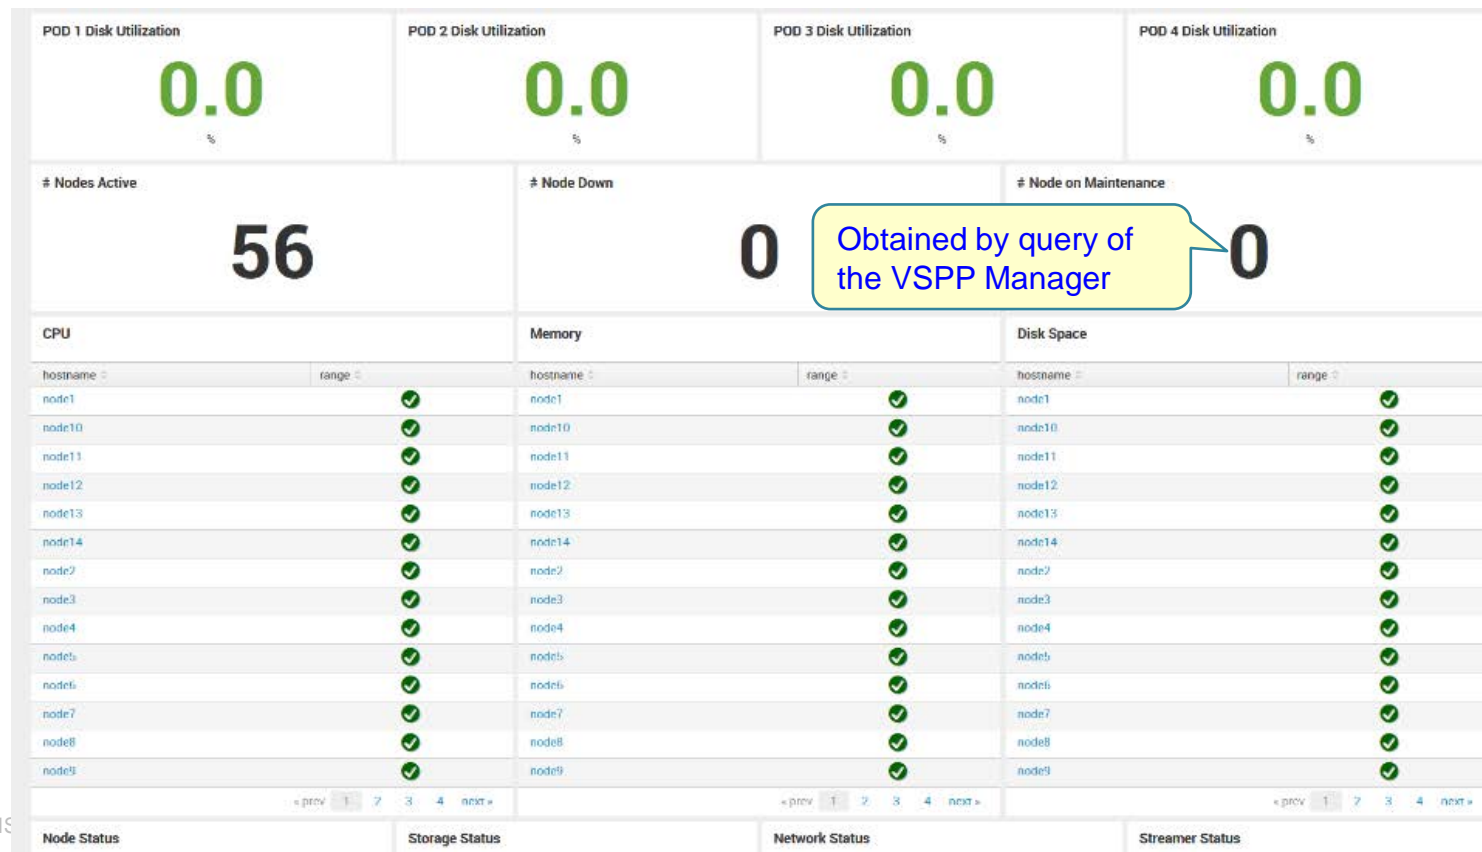


- 5 main system & services operational dashboards
- Real-time monitoring of Level 1 metrics
- Metrics are being updated every 15 minutes
- Selectable time period for Level 2/3 metrics

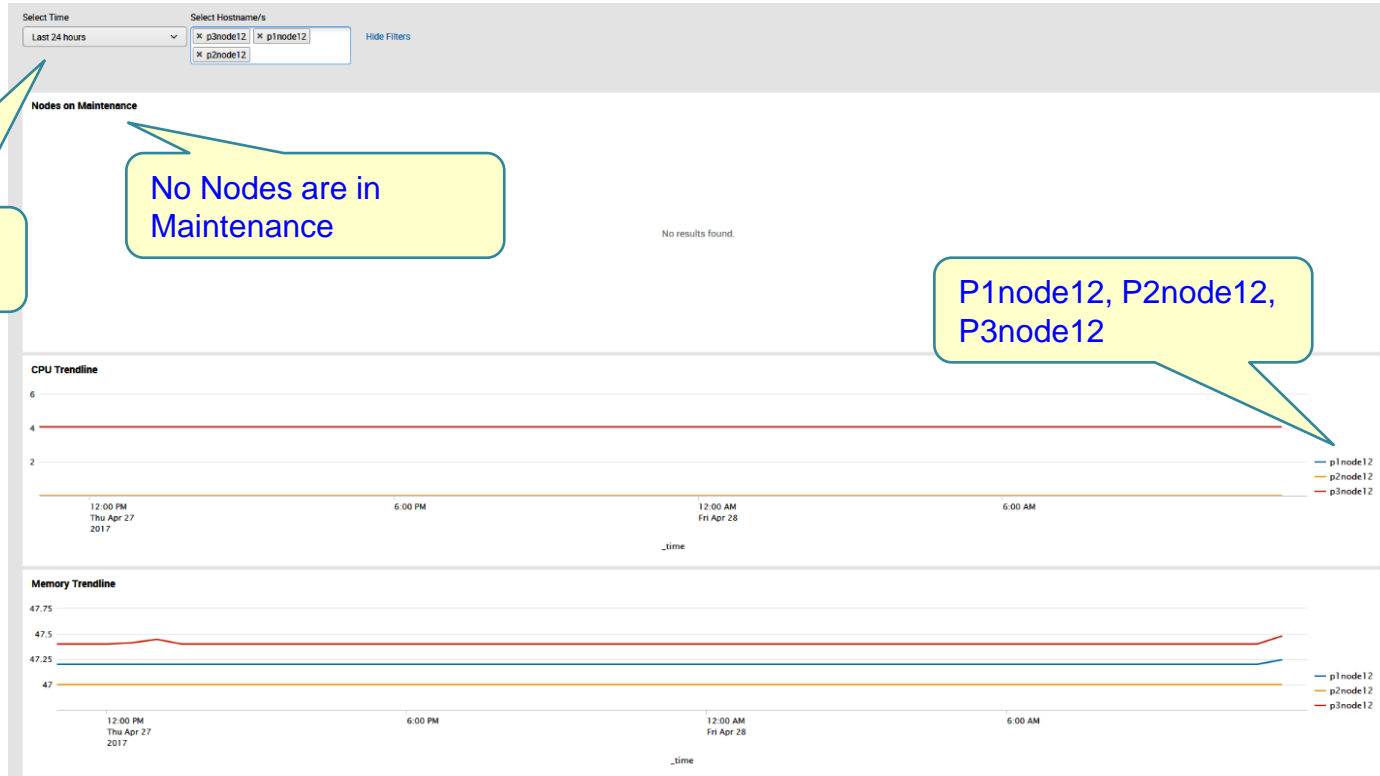
Provide general information on the cDVR system operation such as:

- # of users & daily/weekly user trends for last 7 days/4 weeks
- # of client devices (OVP, Roku)
- # of recorded shows & daily/weekly trends for # of recorded shows
- # of recordings played & daily/weekly trends for # of recordings played
- # of series recordings & daily/weekly trends for # of series recordings
- # of series recording playbacks & daily/weekly trends for # of series recordings playback
- # of playback failures & daily/weekly trends for number of playback failures
- # of recording deleted & daily/weekly trend of deleted recordings
- Top 10 channel watched
- Pod 1,2, 3, 4 disk utilization

VSPP SYSTEM DASHBOARD (LEVEL 2 METRICS)



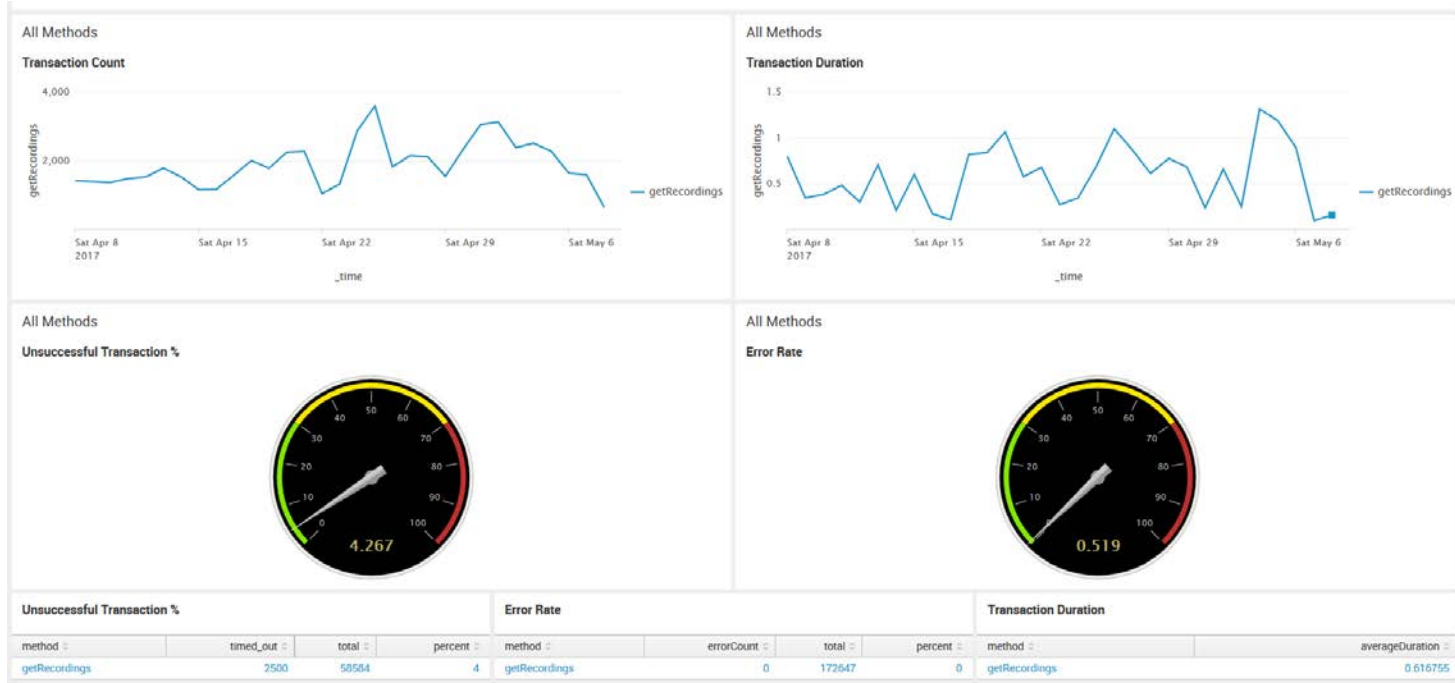
CPU and MEMORY USAGE ON SELECTED NODES IN 24 HOURS (Level 3 metrics)



KUMO getRecordings API STATUS IN LAST 30 DAYS

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% |



SNMP TRAPS & MESSAGES RECEIVED BY cDVR DASHBOARD

| cDVR Component Name | Description | Provided Information |
|------------------------|---|--|
| Network Switches | High-speed Spine and Leaf network switches connecting the VSPP storage and management servers | <ul style="list-style-type: none">• 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. |

Why diagnostics capability is needed?

- When new list of critical or warning issues are generated by the cDVR system, the operational engineer needs to know how to resolve these issues
- Two primary diagnostics capabilities:
 - Metric Diagnostics & Escalation Path
 - Knowledge-based Diagnostics

Metric Diagnostics & Escalation Path

- Initiated when level 1 metrics status is changed to critical or warning
- Suggested list of steps to check based on all the received SNMP traps, GUI messages, and other level 1 metrics
- Escalation path with contact info of Tier 2/3 support engineers for further debugging

DIAGNOSTICS FEATURES

Knowledge-Based Diagnostics

- Some of the observed failures may not be immediately solved by ops. engineering
- When a previously observed failure is resolved, a detailed report is generated with the Best Engineering Practices (BEP)
- Depending on the Root Cause Analysis (RCA) of the observed failures, the threshold-level of one or more level 1 metrics may be changed or new KPIs may be added
- Knowledge-based reports are converted to list of steps for the ops engineers to check before escalating to Tier 2/3 support engineers

- Self-learning capabilities are essential to continuously improve the monitoring system
- Three different types of self-learning capabilities:

Tool Health Check

- Shell script periodically runs and check if one or more background jobs is inactive, and reporting assigned metrics
- Rerun the background jobs that were inactive
- Enable ops. Engineer to monitor the health of the monitoring tool

Self-Optimization (metric threshold-levels)

- Base on historical Splunk logs, and observation of time-dependent patterns of various metrics, monitoring tool automatically reprogram pre-defined initial threshold-levels to avoid false-positive reporting of cDVR system health
- Allow ops. Engineers to reject some of the re-program threshold levels, and enter new levels for specific metrics

Operation Intelligence (OI):

- When storage node starts to exhibit hardware failure or abnormal behavior:
 - One or more storage disk failures (as received by SNMP traps)
 - High-temp inside storage node (iLO alerts)
 - Very high CPU or memory utilization (> 90%)
- Monitoring system sends REST API request to the VSPP Manager to put this node in maintenance mode (for further debugging)
- Possible remedy actions:
 - Replace failed storage disk or other parts
 - Reboot node to check if the node is reporting healthy behavior
 - Perform scheduled software update/upgrade
- When remedy actions are successfully completed, monitoring system sends API request to the VSPP Manager to take the node off maintenance
- Detailed report of the incident is generated

COMPARISON WITH OTHER MONITORING TOOLS

| Monitoring Tool | PROS | CONS |
|------------------------------|---|---|
| SPLUNK> | <ul style="list-style-type: none"> • 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 | <ul style="list-style-type: none"> • Proprietary • Required customer subscription • Potentially expensive |
| Graphite/ Grafana | <ul style="list-style-type: none"> • Open-source tool • Interactive graphs • Scalability • Low-cost subscription | <ul style="list-style-type: none"> • 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 |

COMPARISON WITH OTHER MONITORING TOOLS (continue)

| Monitoring Tool | PROS | CONS |
|-----------------|---|--|
| Nagios XI | <ul style="list-style-type: none"> • Open-source tool • Customized dashboard • Scalability • Server and App metrics • Ability to send e-mail notifications • Generate technical report • Data logs retention and reporting | <ul style="list-style-type: none"> • 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 | <ul style="list-style-type: none"> • 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 | <ul style="list-style-type: none"> • 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. |

SUMMARY

- Real-time Splunk-based cDVR monitoring/alerting system was presented
- Monitoring system dashboard includes
 - Four-level hierarchal server health metrics and KPIs based on pre-defined threshold-levels from all the key cDVR components
 - Server API monitoring and notification
 - E2E report from other cDVR components based on received SNMP traps and GUI messages
 - General cDVR analytics for overall usage and performance
 - Metrics diagnostics and escalation path
 - Self-learning monitoring capabilities, including tool health, self-optimization of threshold-levels, and OI
- Splunk-based monitoring system has many attractive features for high-availability & scalable deployment

SCTE · ISBE

THANK YOU!

Shlomo Ovadia, Ph.D.

Shlomo.ovadia@charter.com

720.279.2875



DENVER, CO
OCTOBER 17-20

