CONNECT
WORLDCAST GROUP

SOLUTION BRIEF

# 10 tips for Successful Monitoring in Media & Broadcast

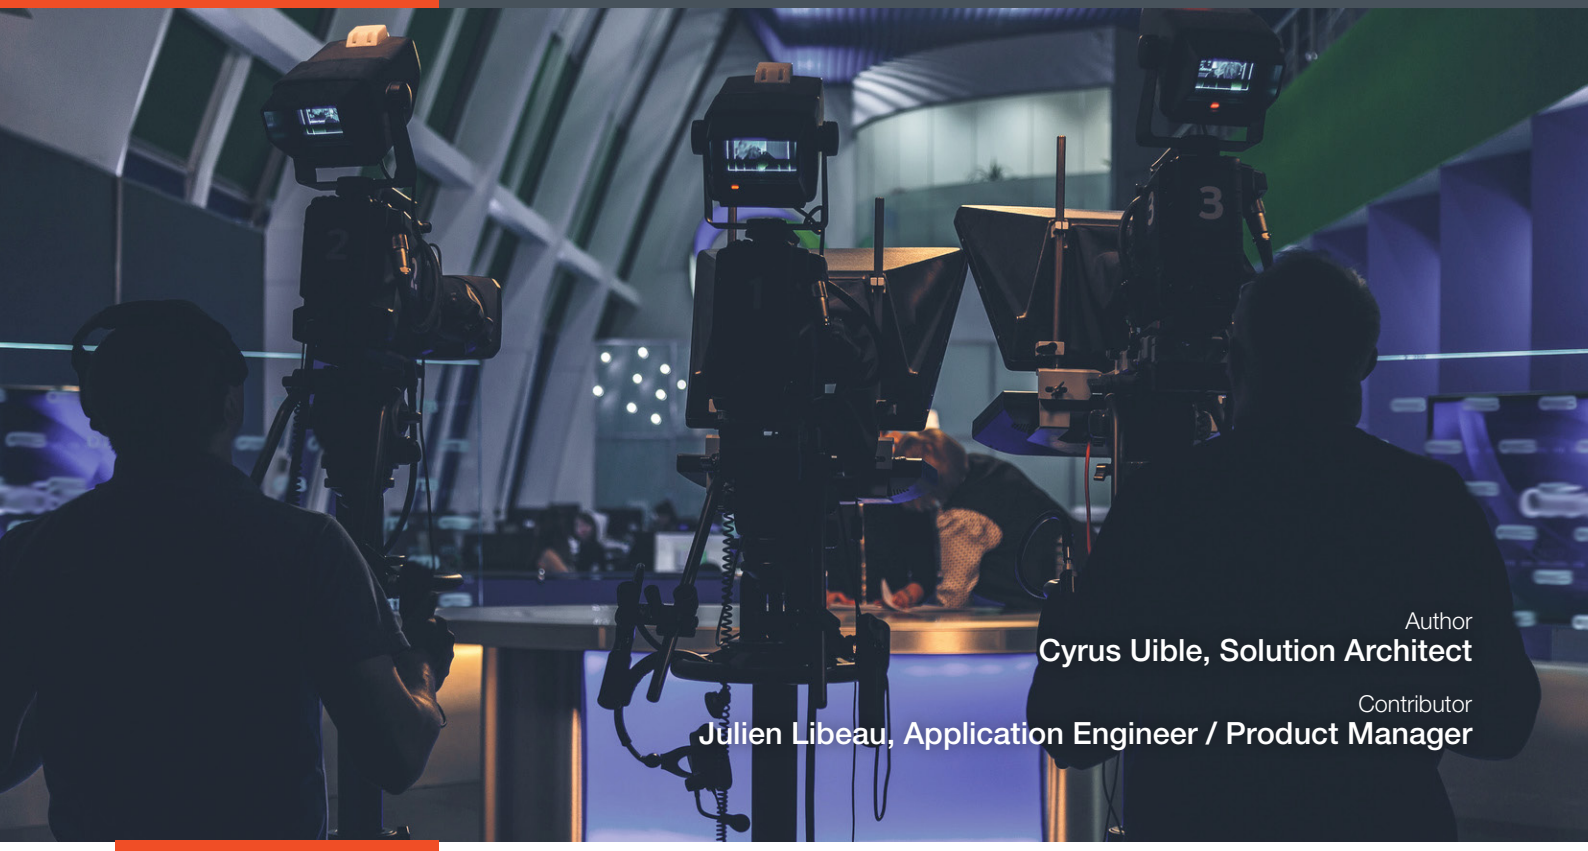

Author
Cyrus Uible, Solution Architect

Contributor
Julien Libeau, Application Engineer / Product Manager



There is no doubt that we live in a world with ever evolving technology. From our mobile phones to cloud computing to virtualization and remote working the things we do today and the tools we do it with are so vastly different than the way we did things a short ten or even five years ago. This is particularly true in the broad scope of the ICT (Information and Communication Technologies) industries.  And even more so when it comes to successful network monitoring and management within these spheres.

We’ve spent the past 10+ years building and implementing enterprise-class network monitoring software for our customers around the world. Listening to their needs and doing our best to execute their vision. But it hasn’t always been smooth sailing. Hindsight is always 20/20 they say and over the years we (and our customers) have learned some lessons that could be helpful to the reader.

CONNECT
WORLDCAST GROUP



TIP# 01

# Planning

Monitoring can no longer be an afterthought at the end of a project. For a few reasons. One is budgeting. I've seen customers plan X amount of $ for a new broadcast facility and after it's built realize that there is no longer budget for a monitoring system. This is a big problem because now they have this fancy new facility with lots of fancy new technology and absolutely no way to know if everything is working without logging in to maybe 10 or 15 different sub-system web interfaces to troubleshoot. Tracking down problems can take many hours (and lost weekends)! The other route is to implement a cheap, perhaps open-source monitoring tool that really can't monitor everything, isn't easy to use, and ends up costing more again in terms of lost time and productivity.

Monitoring can no longer be an afterthought at the end of a project.

But let us say you did budget for a monitoring system. By the way, a good rule of thumb is to set aside 15 - 20% of the cost of all the technology for monitoring of that technology. So, say you are buying $1M worth of tech it isn't unreasonable to plan for $150-$200K for true enterprise class monitoring for that facility. But back to the point. I've seen it happen where a workflow has gaping holes in the monitoring visibility because one of the key sub-systems has no remote monitoring access. No SNMP, no API, etc. This now becomes a new "screen" that must be displayed in the NOC that perhaps was not planned or frankly is not practical because of the number of streams/services.

It's also likely that this particular sub-system has other competitors/alternatives that do have these remote capabilities but monitoring wasn't planned until late in the game. You get where we're going here.

If monitoring is part of the planning from the very beginning, then these major operational issues can be avoided. Keep in mind you don't have to have every single detail. First set aside a rough budget. Then make sure each sub-system in your facility has the key information you need available in some remote way. Don't worry about the delivery. An enterprise-class NMS (such as Kybio) will be able to ingest data no matter the protocol so these device specific details can be worked out later in the project once workflows are better defined. Enough to know right at the beginning that the capability is there.

TIP# 02

# Select the correct architecture

This is specifically related to the architecture of the monitoring itself, not the overall network architecture of either the corporate or broadcast networks. What needs to be decided is where the monitoring server (or servers) will be and what rules/whitelists need to be updated for the monitored equipment to be reachable from those servers.

What also needs to be considered is who will be accessing the management system and from where.
That access needs to be considered as well. Does a centralized model make the most sense? Or perhaps a distributed model where edge devices will do the polling in a number of remote locations. Perhaps there are remote teams of people that each will need to have their own management system with an umbrella layer on top, say at a national super head-end that will have visibility over all the other remote systems.

Tools like Kybio can handle all of these scenarios natively with minimal configuration but there are some licensing and hardware options that will need to be used depending on which architecture is needed.

CONNECT
WORLDCAST GROUP



TIP# 03

# Identify the crucial devices in the network

Outages happen. It is part of the business. The goal of a good monitoring system of course is to spot issues before they become bigger problems. However, that's not always possible. And when the outage happens, the monitoring system has the spotlight. Good or bad, now is the time where all eyes will be on the NMS. Can it tell us what actually happened? Why did it happen? How can we prevent it in the future?

These questions are not going to be answered if the true causes of the outage were not being monitored in the the first place. Or maybe they were being monitored technically but they weren't fully configured to send any alerts. Or maybe alerts were sent, but only to an engineer who doesn't work at the facility any more.

The tip here is to identify every point of failure that can lead to a customer-impacting outage. Think it through carefully from the source to the end consumer. Then make sure these points are not only in the management system but also fully configured to send the correct alerts to the correct people. Or at least create the alarms/tickets needed so the outage can be minimized, or in the best case avoided with automated corrective actions such as switching to backup equipment or chain.

Outages happen. It is part of the business.

TIP# 04

# Configuring your monitored devices

This might be less a tip than a reminder. I've seen plenty of customers do all of the above, only to realize later that they now have hundreds of devices that are timing out in the monitoring system. Why? Because they need to have their remote capabilities enabled through the web GUI before they will respond to any polling requests from the monitoring system. Maybe the API needs to be enabled. Maybe remote capabilities are a licensed option that need to be ordered.

Also, besides just enabling the communication there is further configuration needed. If SNMP is being used there are community strings to set, target management systems to add for the traps. Syslog targets to define. If API's are being used the default credentials should be changed.

These kinds of target device configurations are ideally able to be scripted but sometimes they cannot and it is a painfully tedious process to setup. But there isn't much choice, if you want to monitor these devices it must be done in order to get the most out of your monitoring solution which in the long run will save money and reduce down time.

CONNECT
WORLDCAST GROUP



TIP# 05

# Setting up the alarms

We're getting closer now. We've identified all the critical pieces. Everything has been entered into the monitoring system. Now we need to make sure all the alarms are configured properly and appropriate thresholds have been applied.

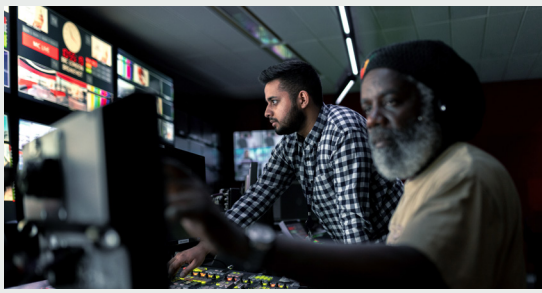

Using Kybio as an example, once a device has been discovered and added to be monitored, some alarms are pre-configured to generate alarms. These should be reviewed by the customer to make sure the threshold and severity of each alarm make sense for that particular customer. My suggestion to my customer is to limit the highest severity (usually this is "Critical") to alarms that can affect their own end consumers. All other alarms can be lower severity depending on what best fits the organization.

The default values for thresholds and monitoring are meant as a convenience and time savor and should be reviewed by the monitoring administrator for their own needs

TIP# 06

# Designing your dashboard

Don't get too caught up here on the term "dashboard". The word can mean different things to different people. My point here is to think about what screen will be displayed and watched by my operators, engineers, maintenance, and management teams? When an alarm occurs will it be clear there is an alarm? Will staff know what to do next?

Some of this involves configuring the screen or visual layer so that it conveys the most information in the simplest way. Some of this may involve training. With Kybio you have the flexibility to show any and all information in the way you want. Whether it be a status board with a matrix of tiles, a signal flow diagram, a rack view of equipment or a floor plan of a remote truck. Any of these are possible - you are only limited by your imagination.

However, the goal should be to not only make it very clear very quickly what the problem is and where it is. All the details should be available with one or two clicks from the main dashboard. Keep in mind here, I've seen dashboards that are very fancy looking but not functional. I've also seen ones that are quite rudimentary but loved by operations because they are simple to use and easy to understand. It's no problem to go for style but by all means do not do it at the expense of utility and ease-of-use

TIP#

07

# Identifying KPI's

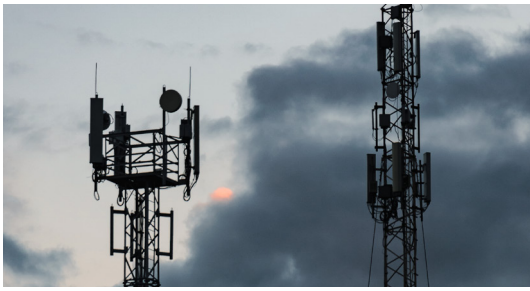

KPI's are your key performance indicators. Now that you've invested both time and money in your management system it's time to get the most out of it. While proper alarming and notifications will help keep you on the air longer and your outages shorter, we also want to be getting better in the longer view. Is our MTTR (Mean time to repair) decreasing over time? Are the number of alarms decreasing? Is customer engagement increasing? These are your KPI's as an organization to help you know how well you're actually performing. Are there blind spots which continue to cause issues for support staff? Is there a particular type of equipment that is responsible for most of the trouble tickets?

Having visibility into these kinds of questions should be possible with a good quality monitoring system such as Kybio. Raw data can be compiled into summary reports and distributed to organizational teams to help everyone stay on track to meet goals.

If your current monitoring solution is not able to help you track these kinds of things it may be time to think of a new system when the next refresh opportunity comes up.

TIP#

08

# Creating Action Plans

The eighth tip we have is to make sure there are clear expectations and workflows in place when issues do occur. In this context, "Action" could be an email to the maintenance team. It could be a flashing alarm on a screen in the NOC. It could be a sms message to the off-duty engineer.

Depending on the context of the alarm it could also mean the creation of a trouble ticket. The execution of a script. The changing of a configuration on a particular device.

Essentially thinking through all the various alarm scenarios that could occur and deciding what actions should follow those alarms. It's a tall task and will take time, but when done properly will dramatically increase the peace of mind for everyone involved.

CONNECT
WORLDCAST GROUP



TIP# 09

# The Monitoring “Lenses”

There are always different ways we look at the world depending on our point of view. The different roles we play can also impact those points of view. For example, my literal view of the world is very different if I’m an astronaut on the ISS than it is if I’m scuba diving on an ocean reef. My concerns and worries are different if I’m climbing a mountain or if I’m caring for a newborn baby. The final two tips are all about different points of view.

In the NMS world, and in this case I’ll use broadcasting/media as a demonstration, I am going to have alarms and status indicators from a large variety of devices and probes. Some of them physical and others might be virtual. Some on-site and others in the cloud.

One “lens” is the system health. What is the health of all my physical hardware that is supporting my workflow.  The fans, the CPUs, the memory usage. This might be best visualized by a rack-view of the hardware.

Another “lens” is the service health. Are the video services running through my data center working as expected.  Are they flowing from the source to the destination with the correct content, correct audio, proper formatting, etc.  This might be best visualized by a simple board of service icons and when you drill down you may see a signal flow for each one.

Another “lens” might be the capacity or resource management. Everything might work ok at the moment but perhaps adding one more service will bring everything to a crashing halt that will mean downtime for multiple services. This might be visualized by pages dedicated to things like network bandwidth or licensing constraints.

There are of course more ways to view things depending on your perspective but these give you a good place to start the discussion with your own teams.

There are always different ways we look at the world depending on our point of view.

CONNECT
WORLDCAST GROUP



TIP# 10

# User Stories

It is important to think about the different roles that are played in your organization when it comes to the design of reports and visual screens. Even decisions as to which devices need to be monitored can be influenced by questions such as “Who needs to know?”

Point number 6 above mentioned this a little but now that we have the different “lenses” mentioned in tip #9 we can imagine how different teams will be interested in these views.

By the way, when I use the term “User Story” I am simply trying to describe how the experience of interacting with the management software will be different for different types of end users.

Let us start with the 24/7 NOC operators. Often, they are the only team that is constantly in front of the management screens. They may be most concerned that the correct programming is being broadcast with the correct audio levels. But they also need to be aware of ANY issues in the system as a whole so they can notify the correct person. It is possible they could be responsible for fixing some things, but for the most point their role is notice problems and notify others that can fix them. So, let’s say they mostly watch the service health page.

The facilities management or IT management might be most interested in the health of the physical devices in their inventory. They want to know if hard drives are about to fill up, if the CPU is being pegged at 100%, if the amount of free RAM is gone, etc. This team might be more interested in the rack view of the monitoring system.

Finally, the technology project team that is planning for the future might be interested in the capacity pages so they can properly budget for let’s say an upcoming major expansion.

All information and presentation then, as you see, has context and should be considered from different points of view to make sure everyone has visibility into the information most important for them.

CONNECT
WORLDCAST GROUP

All together we hope these tips will be useful and instrumental in your next monitoring project and as always feel free to reach out to us for help.

Contact us

HEADQUARTERS
20, av Neil Armstrong
33700 Mérignac Bordeaux-Métropole FRANCE
+33 533 890 500

U.S. REGIONAL OFFICE
19595 NE 10th Avenue Suite A
Miami, FL 33179
USA
+1 305 249 3110

VISIT OUR WEBSITE:
www.worldcastconnect.com