



SEARCH

telecomkh.com

MOBILE TELEPHONY • OPTICAL NETWORKING

RADIOCOMMUNICATIONS • SMART CARDS • TELEVISION • WIRELESS COMMUNICATIONS

you are in: HOME • BUSINESS COMMUNICATIONS • DATA-DRIVEN DEVOPS, THREE KEY TECHNOLOGIES FOR THREE UNIVERSAL OPERATOR GOALS

Data-driven DevOps, three key technologies for three universal operator goals

- TRAINING
- STANDARDS
- SPONSORS
- ASSOCIATIONS

Date: Thu, 10/26/2017 - 15:45

Source: By Pablo Manuel García, Solutions Architect, Blue Telecom Consulting

There are three basic universal goals any telecoms organisation pursues: save costs, make their products more robust and shorten their time to market. It's as simple as it sounds. Everyone wants to produce cheaper, better and faster than the competitors do



Pablo Manuel García, Solutions Architect, Blue Telecom Consulting
Image credited to Blue Telecom Consulting

Seemingly, there is no magical recipe that consolidates them all at once, but we think there's a methodological paradigm capable of getting every organisation on to the right path. By "Data-Driven DevOps" we here understand the convergence of three key technology areas: Monitoring, analytics and DevOps.

Monitoring

Blending into one single concept a few of trending tech topics will make many readers believe that we're just falling for the hype. But let's try to keep our feet on the ground and see how monitoring interconnects that other pieces and closes the circle.

Traditional monitoring is usually based on measuring a set of key performance indicators (KPIs) from the target system, and when an indicator is outside certain predefined thresholds alarms are raised through a messaging system or a dashboard. Those thresholds are usually static and based on the operators' experience. This has two main limitations:

- Decisions of actions to be taken are subjective and depend on the operator. Different people working with the same environments would choose different thresholds for the alarms. We believe it's necessary to objectify this and make thresholds data driven.
- Static thresholds are usually not enough. Not only because all the systems evolve dynamically and so the thresholds would need periodic revision, but because highly seasonal KPIs require more intelligence to effectively detect all anomalies in the system.

Analytics

We're living a data revolution. There's a clear boom in analytics and data scientists are the new rock stars in the labour market.

Big Data tools allow storing and analyzing, in a timely manner, huge amounts of data from virtually any source. In the telecoms sector there are numerous examples where combining environmental data with your own business data this enriches it. This leads to a much deeper understanding of your own business model and can even predict future alarms, so preventive actions can be triggered.

In addition, the latest advances within machine learning and deep learning are currently leveraging a new revolution. Today's algorithms can learn from the past history of your environment and modelise how it should behave under normal conditions. Any change in its expected behavior can raise an alarm to notify something strange is happening.

With enough information, the algorithms may even understand where the anomalies come from, by performing a root cause analysis based on historic data and even identify a solution of the issue.

DevOps

The DevOps paradigm has already demonstrated that it has come to stay. With the quick expansion of virtualisation and software defined technologies, the line between operations and development has become thinner and thinner until it has almost disappeared. IAAS, SDN, NFV and other "software defined things" make configuration files much more important than metal. This revolution leads operations teams to stop writing commands in interactive shells and instead create deployment recipes for their configuration management systems.

In parallel, agile methodologies and continuous delivery processes may also benefit from the integration of an analytics layer. Code versioning, issue tracking and test results can be integrated in the data flow. Analysis will allow identification of weak points in the delivery process and can also optimise the workload allocation between cross functional teams.

From automated to autonomous systems

If you follow the steps described above, you now have configuration methods based on parameterisable recipes as input to your configuration management system. If you were to include a new node in your cluster you would just have to change a couple of parameters in your recipes and then deploy them accordingly.

You also have an intelligent monitoring system able not only to predict future issues but to identify which actions should be taken to solve them. Finally, this monitoring system can directly trigger actions through the configuration management system. Further down, it will get feedback about how the system responds to these configuration modifications applied. This way, it enhances the environment's knowledge about itself by evolving the system's model. This is machine and deep learning put into practice.

Telco Use Cases

With this approach, it is possible to have a solution that autonomously adds new nodes to the cluster when it is expected that its load will grow, let's say due to a programmed marketing campaign. And it can even, via multiobjective optimisation, take into account economic data and improve efficiency through actions such as shutting down nodes when the load decreases in order to optimise the use of resources. This is also highly cost saving in terms of man hours as almost no time needs to be spent on studying a new customer's scenario to properly scale the environment. The product will scale by itself based on its self acquired knowledge about the customer's real needs.

Another use case for software defined networks can be reconfiguring OpenFlow rules to divert traffic in order to avoid certain nodes when it expects a traffic peak in the near future, all this based on social networks talking about its location more than usual (e.g. due to a cultural or social event).

To wrap it all up, an analytics layer on top of data coming from a monitoring solution would allow not only for a better understanding of the environment itself but it could include predictive alarms. When automation enters the equation, it becomes "data-driven", with the following benefits:

- The system becomes more robust as it can act autonomously and in a preventive way.
- Products get to market faster with the use of DevOps methodologies.
- It is more cost-effective as it requires less human intervention both in the deployment and maintenance phases, but also less subjective and less prone to suffer from human errors.

Is your organisation ready to leave the "old school" automation behind and to start embracing Data-Driven DevOps? We believe you and your customers would highly benefit from this, but suggest you test it out in a small project first and take it from there.

print rss compartir

valorar este artículo: ★★★★★

Post new comment

Your name:

E-mail:

Comment: *

The content of this field is kept private and will not be shown publicly.

Homepage:

POST COMMENT PREVIEW COMMENT