NEURO10

Case Study - N10 for System Monitoring

The adoption of digital and cloud-based systems is skyrocketing. Countless physical and virtual servers are being set-up across organisations every single day. Even data centres are growing to be a new industry sector on its own. In 2015, the market size for data centre operators in Southeast Asia was estimated at $1.2bn.

System operators are increasingly feeling the pressure to deliver data and content in larger amounts, at a faster pace and more flexibly than ever before. There is a need for smarter system management solutions, and a big part of system management is system monitoring. Effective system monitoring can not only provide quick reactive responses when something goes wrong, but may also allow preventive measures that would greatly improve service levels.

(1) Machine-driven System Monitoring So That Teams Can Do More with Less

> **Scenario**: The IT team in CompanyX is reponsible for ensuring the performance and uptime of the IT systems. These system operators mainly monitor systems by setting threshold rules on critical system metrics that trigger an alarm when any rule is broken. Example of a rule: If CPU usage of serverA exceeds 80%, trigger a "Warning" alarm.
>
> However, the growing number of servers and threshold rules in today's large and complex IT systems makes system monitoring overwhelming and much harder to manage. For example, as systems evolve over time, system operators have to continuously add, edit and delete rules. This is challenging not only because it is tedious, but also because of the multi-fold complexity as there is dependency across servers and rules to consider.

Solution: N10's anomaly detection module (ADM) learns the norms and deviations in system behaviour over time via data and automatically triggers an alarm when an anomaly is detected. This model can also adapt independently at real-time to changed environments.

To enable ADM for any particular time-series data, include the following short block of code in the data-forwarding script -

```
pingData = pingData.toString();
pingData = pingData.replace(/\{|\}/g, '').trim();
pingData = pingData.replace(/\s/g, '').split(",");

var measurement = {
        'instance' : ipaddr,
        'loss' : {
                'value' : -1,
                'unit' : Constants.UNIT_PERCENT,
                'alarm' : Constants.ALARM_CLEARED
        },
        'rtt' : {
                'value' : -1,
                'unit' : Constants.UNIT_MS,
                'alarm' : Constants.ALARM_MACHINE_LEARNING
        }
};
```

N10's ADM will observe this time-series data and learn patterns over time. ADM then predicts future values based on the learned system behaviour and regard them as normal values. When the actual value is eventually received, ADM compares it to the predicted norm and if it deviates, an alarm will be triggered (note: the extent of deviation may be configured by the user, i.e. user may configure ADM to be more/less sensitive which triggers alarms with less/more tolerance to deviations) -

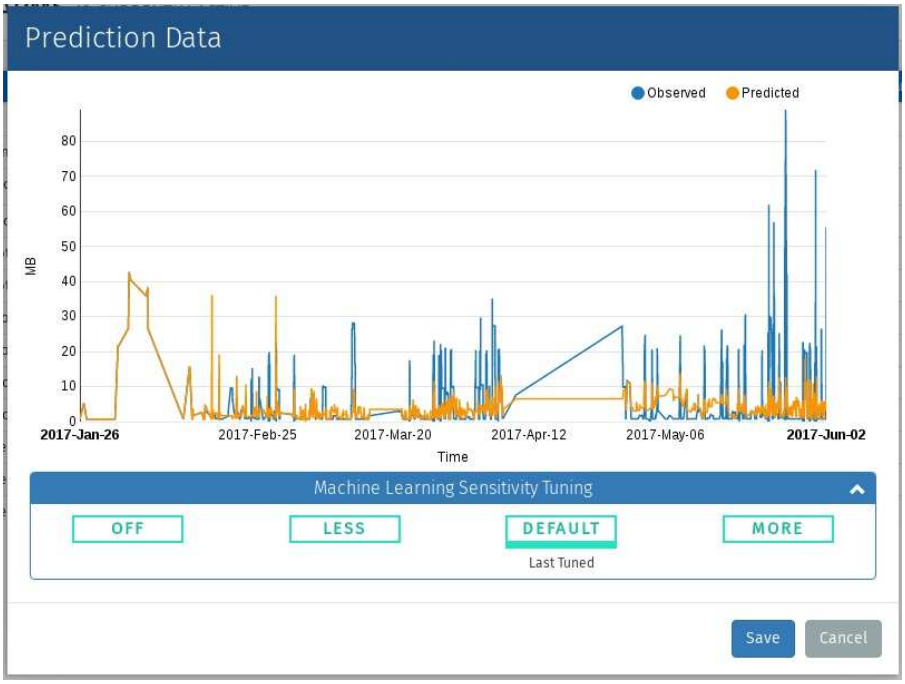

ADM helps the IT team move from rule-based monitoring (static) to model-based monitoring (dynamic). Rules are managed manually, while models evolve with data. With ADM, the IT team needs to worry less about managing rules and can instead focus on servicing alarms, which improves mean-time-to-repair (MTTR). The move from reactive to predictive monitoring may also help to reduce mean-time-to-failure (MTTF) since warning signs may be picked up early before they lead to major failures.

(2) All-in-One Alarm Management Interface for Quick and Comprehensive System Monitoring

> **Scenario**: As business processes and IT systems of CompanyX evolve over time, so would the behavior of alarms. For example, some system components are more actively used in certain times of the year (hence higher occurrence of alarms), but their activity drops in other times of the year. The IT team needs to build fast and good understanding of how the alarms are behaving to update their work plan (e.g. SOPs) so that they can manage these alarms more effectively.

Solution: N10's web platform aggregates and analyzes all alarms on a single interface for an instant overview of alarm behavior and trends.

To view the list of alarms and organize them by filter or sort -

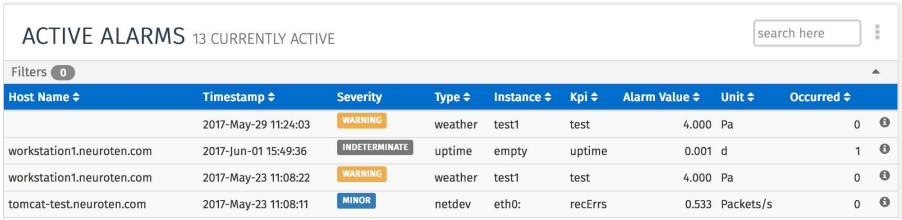

To understand the activity of alarms across different days as well as other trend analyses -

From the above charts, the IT team could quickly observe that large alarm waves tend to occur every 3-4 days. The team may plan to increase manpower availability to service these alarms at such time intervals.