REPORTING ON MONITORING DATA

GroundWork White Paper by Thomas Stocking

Introduction

This white paper describes the ideas we have used to assemble the reporting systems in GroundWork Monitor. This article focuses on the concepts that make great reports possible, and will hopefully give you some ideas of what's important for you to report on, and how. For specific product documentation and details, please see GroundWork Support.

What reports do you need from a monitoring system?

Reports are the main way organizations communicate Key Performance Indicators (KPIs) for the business to management staff.

Monitoring systems gather a lot of potentially useful KPIs, but the primary purpose of monitoring systems isn't actually reporting. Instead, it is primarily alerting DevOps when things are bad, or about to get bad. It's also providing searchable data to DevOps so they can diagnose, troubleshoot, and tune systems.

Reporting is a secondary function, and so it is sometimes ignored in the early stages of monitoring deployments. As deployments mature, however, reporting becomes a mainstay of capacity planning, forecasting, and many other data-driven business processes. Monitoring systems have a part to play.

Reports vs. Dashboards

Do you even need reports? Sometimes, reports are overkill.

Often all we need is a dashboard with some historical data in context to make the same inferences and gain the same insights that a report would show. Also, reports sometimes bury *information* in mere *data*, and dashboards can bubble up relevant information to the top of a colorful, interactive display in ways that make the insights you really need clear.

Reports on data, therefore, show their value over time, more than on a short-term basis. Looking at monthly or annual summary reports can help put dashboard data in perspective.

Dashboards also perform according to the inverse of the amount of data they present - the more data, the slower the dashboard. If you find that you keep adding data to your dashboards until they are overloaded, you probably need a report instead.

SLA Reports

Having Service Level Agreements (SLAs) for critical business services is a great way to manage costs, expectations, and business continuity planning.

It's possible to have systems that are *never* down, but these systems can cost far more than systems having 99% or 95% uptime. An SLA can build in cost estimates, staffing budgets, etc., which can put these investments in perspective.

Monitoring systems can track compliance with SLAs. Indeed, this is often the driving reason to implement reporting on the monitoring system, or to have certain monitoring profiles in place. SLA reports need specific features to remain relevant, however.

Business Service Monitoring

A simple rules-based aggregation of individual monitoring metrics can deliver the overall state of a business service as a single object. This is what Business Service Monitoring (BSM) does. Using BSM can greatly simplify the crafting of SLAs.

Re-processing of Events

In cases where a monitoring system *misses* a business service outage, it can be necessary to insert downtime into the SLA computation after the fact. While in an ideal world, all business services are perfectly monitored, the reality is that unforeseen outages of some systems can occlude monitoring visibility. It's also true that not everything is always monitored as it should be, and monitoring systems need to be extended to cover new or changed infrastructure. While this can and should be largely automated, the lag in doing so can at least temporarily affect visibility. This must be accommodated by the SLA reporting system.

Re-processing events can also help eliminate *false-positive* monitoring, where a rule or threshold was improperly set to be too sensitive. In this case, false outages can be removed after the fact, again, making SLA calculations more accurate.

Event Reporting

One of the global KPIs of DevOps is *number of events*. Simple reports that track the count of events, by type, for each major unit of infrastructure can be very useful in ensuring service continuity.

For example, disk errors on a storage array are typically logged and can be detected by the monitoring system. Insignificant in themselves, in aggregate they can indicate issues with array controllers, or possibly specific storage units. Units for which reports indicate more errors may be moved up in the maintenance schedule. Disks themselves also vary in quality, and error events can be indicative of a bad batch of disks. Reports on these errors can then be used to negotiate with suppliers, etc.

The format of event reports need not be too specific. It should be possible to determine from the report what the relative counts of events across groups of monitored units are, and allow for the precise determination of the time of each event in the counts.

Reports and Tuning

While dashboards are useful in determining the current state of KPIs in an immediate historical context, reports can often reveal longer-term trends and changes that dashboards will not. Tuning of systems for optimal use of resources, therefore, is sometimes easier to achieve when digesting information that has a longer time scale, i.e., reports.

Performance data metrics from systems infrastructure is best represented graphically, rather than in a table. Liberal use of graphics can add significant value to such reports. Creating charts and graphs spanning longer time scales can be processor intensive, so some accommodation for this may be needed when planning how and when to generate reports involving them.

Ideally, reports of performance data should include at least basic statistics such as average, max, and min values. Variance, correlation, trending and predictive analytics are a plus.

Staff Management and Operations Reporting

Operations staff are often managed based on productivity measures, such as tickets closed or hours worked. While infrastructure monitoring systems are not comprehensive in terms of measuring productivity, certain metrics may be available that will be of interest for apportioning credit to teams or even individuals.

Numbers of notifications to staff or teams in a given month can be used as a proxy for workload, and outage time (especially outages related to SLA relevant events) can be regarded as a measure of efficiency in its inverse, that is, small number of short outages mean an efficient team. These numbers can often be determined through efficiency reports generated by the monitoring system.

How are reports generated?

Ok, so now we know why we might need reports from the monitoring system, and what to look for. How exactly these reports are composed, generated, and distributed will vary significantly from organization to organization, with some requirements being very detailed, specific, and regular, some more general, long-

term, and ad-hoc. In all cases, however, you should look for ways to generate reports that minimize impact on business operations, and maximize the utility of the reports themselves.

Automating Reporting

Reporting is a chore, and is often delegated to junior staff or outside experts. This is not necessarily a mistake, but the true value of reports may not be realized unless some investment is made to both design them properly and to automate their generation.

Some reports, such as SLA reports, are fairly simple in design. The complexity comes in the crafting of the SLA, deciding what to monitor, and how to maintain that monitoring over time.

In the case of custom reports, it's important to capture the value of investments made in up-front report design, and not re-do the same work each week, month and year. A little up-front work by experienced staff to properly specify the reports needed will likely pay off, while less specialized staff can be efficiently tasked with the technical work of setting up their regular generation.

Use of reporting tools with scheduling features, scriptable interfaces, and exportable formats (XML, CSV, PDF) will facilitate the automation process.

Targeted Reports for Incidents

Sometimes a report has a specific goal in mind, such as listing all events similar to a specific recent one, or comparing the configuration of multiple systems or settings. Often a report can be included in the incident response process.

While every incident is unique, and modern IT systems tend to be extremely dynamic, there are some parameters of incidents that are common, such as:

- When the incident started
- When and if revenue generation was affected
- When and if customers were impacted
- When the incident was detected
- When the incident recovery started
- When the incident was resolved

The monitoring system can potentially provide all of the above parameters in a report, if properly instrumented and designed.

Distributing reports

There is no utility in a report no one reads, so generating reports should be done selectively according to the principles above. However, once the reports of value are generated, they have to be delivered to the person

or people who can use them to make decisions. There are several things to look for when deciding how to distribute reports.

PDFs

Reports should be in a form that is set, so everyone is looking at the same data. Editable reports, with the exception of the reprocessing of events for SLA accuracy mentioned above, are not really a good thing. The PDF format is a good choice. Not that you can't edit a PDF (you can), but it is designed to be a read-only format for most applications. Your report generator should support the creation of PDF versions. You might also consider HTML as an output format. There are advantages when distributing reports as HTML pages.

Emails

Emailing PDFs is one way to distribute reports. Beware of the drawbacks, however. Emails (especially automatically generated emails with attachments) can be filtered or ignored. There are a lot of emails to process for a typical executive, and yet another recurring message with a report to read can be forgotten easily. Also, attachments add processing time and storage space to email systems. PDFs are not typically very large files, but some reports, especially event reports with full detail, can be quite long. If you pay for storage on your email system, make sure you're not adding an ineffective cost by attaching reports in email.

Scheduled Deposition

Once a report is generated, it can be placed on a shared storage area and accessed when needed or wanted. You can use any of a number of places for this, such as a company Dropbox, SharePoint server, or even an *htdocs* directory on a commonly accessible internal web server. This method allows you to always have access to the reports you need, when you need them, centrally. Old reports can be backed up, archived, aged out, and preserved for compliance or policy reasons, all from one place. PDFs can easily be linked to in this way, but are not as dynamic as HTML. If you do decide to deposit the generated reports in a central place, consider HTML as an output format.

Managing data

Reports are based on data. Data is used at several levels in monitoring systems, for determination of state, for graphing, and for tracking the time and duration of events, scheduled downtimes, and more. Not all the data is useful for reporting, and in fact, reporting can interfere with the other uses of the data itself.

Production vs Archive

One important factor in maintaining a monitoring system is the management of production databases. It is important to preserve resources like storage, memory and CPU for running queries used in determining state, presenting dashboards, etc. Reports by contrast are typically long-running jobs, taking from several minutes to

hours, potentially, and can use a lot of CPU. There's no need for a report to be real-time; it is designed to be run after the fact.

Also, many reports can be prepared against a subset of the data used in a full database. Archive copies of production data should be made each day, incrementally, and reports should be made against these archives instead of the dynamic, real-time versions. The production database can be trimmed to optimal size and maintained, reindexed, etc after the archive process to keep it in optimal condition.

Data Retention

While each organization will have its own requirements for data retention, there are some principles we have found apply to retention when considering reporting.

Often reports are useful as comparative measures, day-to-day, week-to-week, month-to-month, and year-to-year. Thus, it makes sense to retain at least 13 months (we use 56 weeks) of data to allow a single report to show summary stats for this month, last year. In GroundWork Monitor, we actually also use this in some dashboards, but the principle is the same. Archive data shouldn't grow forever, either, but it can of course get much larger than the production data can efficiently grow to. A few years are typically practical to keep on disk, and so your organization's data retention policy should be possible to accommodate with an archive database..

Conclusion

We hope this description of the issues and ideas we have encountered in reporting on monitoring data is useful for you in creating a reporting solution for your organization. If you are interested in crafting SLAs, running the pre-defined SLA reports, and creating your own custom reports on monitoring data, we hope you will explore the tools and documentation for these tasks using GroundWork Monitor Enterprise.

Have questions?
Email us at support@gwos.com.