FINDING THE NEEDLE IN THE HAYSTACK

retrospective

Index

RETROSPECTIVE - Whether you develop Software, integrate systems, have to carry out massive tests or are in support and struggle with problematic tickets, Retrospective enables you to instantly search through distributed logs and get instant data transparency. In addition, if you need to keep a «real-time» eye on different logs, Retrospective offers you a «tail» feature that conveniently merges the different tailed logs into one unified view.

retrospective

Chapter 1.
Bumper harvests

According to the fifth annual IDC Digital Universe study of 2011[1] the world's information is doubling every two years and in 2011 would reach 1.8 zettabytes. It then goes on to predict that by 2020 the world will generate 50 times more information with 75 times the number of «information containers».

That means an awful lot of additional hay to be added to the already mountainous stacks. This can only aggravate the current data management crisis that is facing all companies, great and small; How to separate the wheat from the chaff.

Originally intended as a tool for diagnosing and debugging code, system logs evolved into a way of recording transactional and event information and then became increasingly used for system troubleshooting, forensics and security incident response. With the increase in legal requirements for compliance with audit or security policies, the size and number of logs exploded.

To add insult to injury, the current expansion in information production, transport and consumption has led to increased network activity and the rapid expansion of systems and infrastructure. More events happening across more devices means storing even larger amounts of data in even more log files. Systems administrators and analysts, already struggling to effectively search a wide variety of log files, let alone extract actionable information from them, will need the help of good log management and analysis tools more than ever. Even as the amount of data to be analyzed increases, businesses are finding it useful for a growing number of purposes.

This is borne out by the recent SANS Sixth Annual Log Management Survey Report[2] which states that one of the major log management issues today is «he ability of log management systems to deliver value from the logs being collected, specifically in the areas of searching (where 36 percent of respondents reported problems), and analysis (where 34 percent had problems)».

[1] IDC 2011 Digital Universe Study: Extracting Value from Chaos
[2] SANS Sixth Annual Log Management Survey Report 2010

Chapter 2.
Bringing in the sheaves

So what features should a self-respecting log management tool provide? Irrespective of whether you keep logs for compliance reasons, security incident responses, forensics and audits, or to help maintain your systems health and improve it´s performance, you have three basic tasks:

1. Identify and access the relevant log files anywhere they are located, in any format.

2. Query the logs and extract pertinent data using search queries and filters.

3. Turn the extracted data into actionable information and present it effectively

Log management products should be judged by the degree to which they assist you in carrying out these tasks efficiently. The good ones will provide, out-of-the-box, parsers for the more common log file formats, indexing and regular expressions functions for searching, templates for standard compliance reports, and event correlation for automatic detection and notification of critical events.

So the real differentation between the currently available products is how well they implement the more low-level funcionality, such as the data aggregation methods used, the combination of matching and parsing for searching, how well they provide for data normalization and correlation, and how good are the management consoles. And last, but not least, how easy it is to deploy and maintain the product.

retrospective

Chapter 3.
What's out in the field?

The current market offerings seem to fall into two main flavours of Log management products, all-inclusive appliances or software-based products. The all-inclusive appliances tend to be limited to a few standard configurations and are usually the outside the scope of this paper.

The software-based products range from standalone applications to full-blown SIEMs (System Information and Event Management). SIEMs however can be expensive, depending on the size of the network, and are aimed mainly at large companies with extensive networks to protect and maintain. They are also difficult to deploy, due to the need to distribute, install, and configure additional software across multiple clients, and require dedicated and trained staff to manage it.

The standalone applications are the agile answer to the bumper harvests in log data. They are easily deployed, easily updated and have minimum system-impact. This makes them particularly useful to those focussed on i ncident and problem analysis, enabling them to progress further and faster with incident analysis before having to escalate it up the management chain.

retrospective

Chapter 4.
Winnowing the wheat

Retrospective is a good example of a standalone application. Retrospective is a tool to help you analyse individual incidents and problems occurring in your integration layer and correlate this information in a meaningful way to assist Incident and Problem analysis according to the ITIL framework for identifying, planning, delivering and supporting IT services to the business. It does this by providing the following features to address the main concerns of the data analyst.

Deployment

Retrospective is a stand-alone, small , non-invasive software package with an Easy Configuration wizard to help you set up profiles and add your data sources (devices and files).

Search

Search performance has been improved by implementing the indexing of log file meta-data, extending the search filters to include, testing of the string to search, data time and Boolean operators. Retrospective is now also able to save and load search filters.

Analysis and Presentation

Retrospective has improved its analysis and presentation features, implementing «tail» functionality to provide real-time viewing of incoming data and show major trends, as well as providing support for Log4j package filtering.

The application also provides the ability to present differently structured files in one view and to define a column split for structured files.

All of this is managed via a Google Chrome-like interface with expanded tab functionality, allowing you to drag and drop tabs between windows or undock them completely. You can also explode all the tabs to separate windows or implode them back into one. In addition, Retrospective allows the saving and export of the complete desktop.

This makes Retrospective the perfect, cost-effective tool to enablie the specialists to start analysing a particular incident or problem, often with very basic information such as a single reference number. This key information can then be used to extract further relevant information to trace what actually happened in the system and present that information in a way that will aid the swift resolution of incidents and problems.

retrospective

Chapter 5.
The final straw

«The amount of log data being collected is growing at the rate of 15 to 20 percent per year,....Some companies report a 100 percent increase per year. The top factors that respondents expect to impact future growth are 'Increased log sources'».[3]

It seems that the current problems companies are experiencing analyzing data from all the different types of logging devices and extracting all the value those logs have to offer.can only get worse. It is of no use to able to get all the data you need if you have to work too hard to extract it and then make sense of it.

The industry consensus seems that immatureter searching and reporting capabilities are to blame for this situation. These issues can only be addressed by improving the reporting features provided by log management products. This can be achieved by enhancing the features for correlation, presentation and decision support capabilities with added intelligence in the analysis process.

In addition, it seems that many businesses now want to see improvement in the use of log data to enhance day-to-day operations and are more aware of the benefits that efficient access to logs can bring by helping reduce the cost of problem resolution and thesupport of other operational needs.

The log management industry is obviously entering its mature phase, with the major concerns now shifting from log data collection and storage to focusing on how much businesse can do with all that data their logging systems are collecting. So those log management and analysis vendors who want to survive and prosper, take note.

[3] SANS Sixth Annual Log Management Survey Report 2010

retrospective

Chapter 6.
Recommended reading

Extracting Value from Chaos June 2011, John Gantz and David Reinsel
available from: http://netherlands.emc.com/collateral/analyst-reports/idc-extracting-value-from-chaos-ar.pdf

SANS Sixth Annual Log Management Survey Report April 2010, Jerry Shenk
available from: http://www.sans.org/reading_room/analysts_program/logmgtsurvey-2010.pdf

Search log analysis: What it is, what's been done, how to do it 2006, Bernard J. Jansen
available from:http://www.sciencedirect.com/science/article/pii/S0740818806000673

Choosing Your Log Management Approach February 2008, Anton Chuvakin
available from: http://www.slideshare.net/anton_chuvakin/choosing-your-log-management-approach-buy-build-or-outsource

InfoWorld review: Better network security, compliance with log management August 2010, Roger A. Grimes
available from:http://www.infoworld.com/d/data-explosion/infoworld-review-meeting-the-network-security-and-compliance-challenge-658