Global Information Assurance Certification Paper

Log Analysis as an OLAP Application
- A Cube to Rule Them All -

Submitted by Leong Ying Siong Clement as practical
assignment (v1.4b Option 1) for GIAC GSEC certification

Abstract

Log analysis is an integral part of effective security management. The logs that an organization keeps are only as useful as the rigor of analysis that they are subject to. That rigor would have to stand up to both the voluminous amount of logs as well as the myriad of logs within an organization that carry security information. Fortunately, another problem of similar veins has been solved years before in the business world by the means of OLAP (Online Analytical Processing) and the technology has been featured in several log analysis products recently. However, most of such tools are either rigid in their queries, or are pricey and many potentially lock the customer to proprietary solutions.

This paper discusses a specific implementation of using OLAP technology on log analysis, in particular by using the Seagate Analysis OLAP client. The Seagate Analysis OLAP client, which is released free to registered users since February 2000, snuggly fits into this role for log analysis. This tool is free and powerful enough to be the first step for practitioners to explore OLAP's utility. We will discuss how OLAP alleviates the log analysis problem, basic ideas on OLAP and related database design concepts. There is also an iteration through a mini project that uses the Seagate Analysis on Windows NT Event Logs.

Overview: OLAP and Logs

"If a tree fell in the forest and no one is present, did it ever fall?"

The merits of log analysis in providing a security overview of your systems are well articulated elsewhere. I chose to quote the ancient philosophical question to epitomize the point. Interestingly, if the tree did fell, it becomes a log. Anyhow, there are many tools in the area of log analysis. A tour around Google would review Counterpane, SANS and SecurityFocus to host much insightful contents in this area.

What is apparent from this tour is that logs are getting tackier to handle, having grown exponentially in both volume and variety. Databases are now called in to manage the unwieldy logs. Respectable log management products would now come bundled with a log analysis console reporting on canned aggregates: top ten visitors, top ten intruders, top ten most popular pages, top ten most visited weekdays. However, for those who have built their logging infrastructure on open source or free tools such as Snort, MySQL, Syslog and home-brewed Perl scripts, we would be hard pressed to find something equally free and powerful to fill the role of log analysis. The Seagate Analysis (SA) OLAP client fit this role nicely.

The benefits of using SA on your horde of logs are as follows:

- If you are not already using a database to hold your logs, you probably should and reap the benefits that Zbyszek Sobiecki articulated in his 2001 article to DaemonNews. You'll get more mileage from your logs when they are in tables.
- The power of SQL in finding nuggets of information from a database of logs is briefly mentioned in Anton Chuvakin article to SecurityFoucs in 2002. However, as powerful as SQL might be, it is not a straightforward language to use if complicated queries are to be issued. The SQL statements that are formed may not capture the essence of the question that the analyst has. SA provides a user interface that nicely mitigates this problem.
- Log analysis tools that uses database are often limited to a set of often-used queries rather than allowing the user to define particular queries which they are interested in. SA as a dedicated OLAP client exposes more functionalities to the user's control.
- SA is part of the popular Crystal Report suite of software and has a limited but still useful reporting engine built in. SA also builds colorful graphs in 3D.
- SA can use data sources other than databases to access data for analysis. For example, SA can use the NT Event Logs, Microsoft IIS logs and Lotus Notes logs.
- SA does not handicap SQL wizards. SA allows for direct SQL statement formulation against the data sources and assists by providing helpful interfaces.
- SA brings the power of Online Analytical Processing (OLAP) to your logs. OLAP is used for decades in the business world analyzing business data for both strategic and tactical insights. A particular trait of OLAP called 'slice & dice' will be explored in this paper. SA brings new level of accessibility to various aspects of your logs via this feature.
- Finally, SA has been offered free by Seagate Software since 2000 in a bid to showcase their product. It costs you nothing to try and even keep it, if you can find it. However, that situation changed in 2002. More on that at the end of this paper.

If you do not already own a copy of SA, you may find its limited availability on the web somewhat of a letdown. Nonetheless, as the main subject of this paper is on the

application of OLAP to the area of security log analysis, the concepts discussed here would be equally applicable to other OLAP clients. In any case, to effectively use SA or any other OLAP clients, you would need to understand the basic tenets of OLAP.

Basic OLAP Tenets

The basics of OLAP are also well documented elsewhere. I will however provide a short treatise here for those eager to jump in to use SA on their logs. This section will not do justice to the deep expertise developed in this area over the years; the reader is encouraged to peruse the references listed in the reference section for further readings.

OLAP describes a set of technologies that allows analysts to quickly gain answers to the 'who' and 'what' questions premised on a, usually large, set of data. OLAP applications typically achieve this through multidimensional views of aggregate data derived from the data set. OLAP also answers tougher questions such as 'what if' and 'why', but this will not be the emphasis of this paper.

In an OLAP application, there are dimensions and measures. Dimensions represent the objects of analysis. For example in a network intrusion log, dimensions could be source IP address, destination IP address, attack signature and time. Measures represent the numeric data that you analyze across dimensions. For example, in the same network intrusion log, the number of instances of connections between a particular pair of IP addresses would be a measure. So would be the following arising from the dimensions we cited:

- How many of a particular attack is seen to be attacking or originating from a particular IP address
- How many particular attacks are common over the day, week, month or year
- Which particular IP addresses are attractive to attackers over time

The slice & dice feature of OLAP makes these queries intuitive and easy to use. We will illustrate more on this feature in the following example using NT Event Logs.

Dimensions are also hierarchical. For example, attack signatures can be grouped into categories like denial of service, privilege escalation or just probes. Measures can then be calculated according to such groups that can be broken down into specific events within each category. This 'drill-down' capability allows examination of data at various levels depending on the analysis at hand. However, most of the hierarchical relationships that are natively defined are the time dimension having children dimensions such as year, quarter, month, day and even hours and minutes of the day. Other hierarchies would need to be explicitly defined for each data set.

The dimensions and measures forms the data cube that the analyst manipulates to access various views of the data. The following diagram illustrates 2 such cubes from the above

examples:

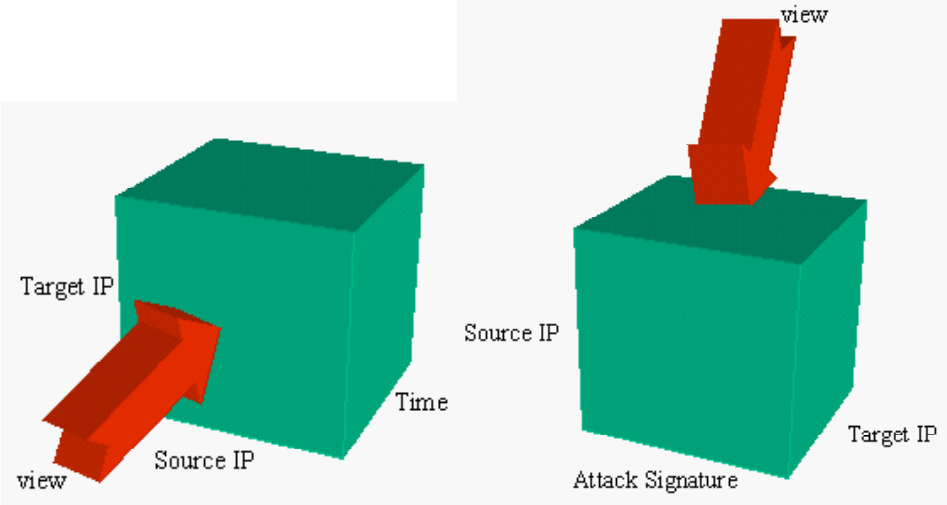

Cube A Cube B

The cubes differ only in the dimensions that are used to build them and the red arrows represent views that the analyst is using to examine the cube.

In cube A, which is formed by the Target IP, Source IP and Time dimensions, the analyst will see the total number of connections (i.e. the measure) between target IP address and their respective source IP addresses over all time that the logs has been kept for. The slice & dice feature of OLAP tool allows the analyst to slice the cube along the Time dimension and examine the same measures over a particular month or quarter.

In cube B, which is formed by the Source IP, Attack Signature and Target IP, the analyst will see the total number of attacks types used on each of the target IP address over all source IP address (i.e. irregardless of the attackers). Here, the slice & dice feature of OLAP will allow the analyst to slice the Source IP dimension and examine a particular IP address' attacks on the target machines.

The power that OLAP brings to a database full of logs should be quite apparent from these examples. OLAP's utility in log analysis can be only leveraged further by the richness of the logs and the variations of queries that the analyst can conjecture. Raw logs can be augmented with ancillary information to surface further information through OLAP. For example, IP addresses can be resolved to either their domains (i.e. .com, .gov, or .edu) or geographical assignments (i.e. assigned to USA, UK or China) and be subjected as additional dimensions for analysis.

Using Seagate Analysis (SA)

SA has a easy to read and illustrative help feature to guide the new user. The publisher Osborne, which publishes "Crystal Report – The Complete Reference" also used to have

the chapter on SA freely available on the Web in PDF format but may have been withdrawn from the Web but will certainly still be found in your favorite bookstore in a book on Crystal Reports. We will go through the essential steps in using SA but SA has so many features that it is worthwhile perusing either of these documentations or many others that you would be able to find on the Web or your nearest bookstore.

Download and Installation

Though SA is released free to the public in 2000, it has proved to be a hard to find download recently. Crystal Decisions' support for the free version of the program is also limited to a knowledge base and a forum specific to SA. Nonetheless, SA binaries can still be found in a corner of the web via Google. Installing SA is as typical as other installer-assisted application and we assume the user would have no problems installing SA.

Taking SA out for a spin

The quickest way to test drive SA would be to train SA on the NT Event Logs. We will use this as an example to illustrate SA's capabilities in analysis.

We will assume that the SA is installed on the NT or Windows 2000 machine that has event logs to be analyzed. Upon starting SA, it will ask whether to start a new analysis document using the Query Assistant. The Query Assistant is a wizard to guide you in creating an analysis document and would be helpful to use when you are exploring other aspects of SA.

If you have not installed and executed SA so far, the following screenshot will be the first time you come face to face with SA.

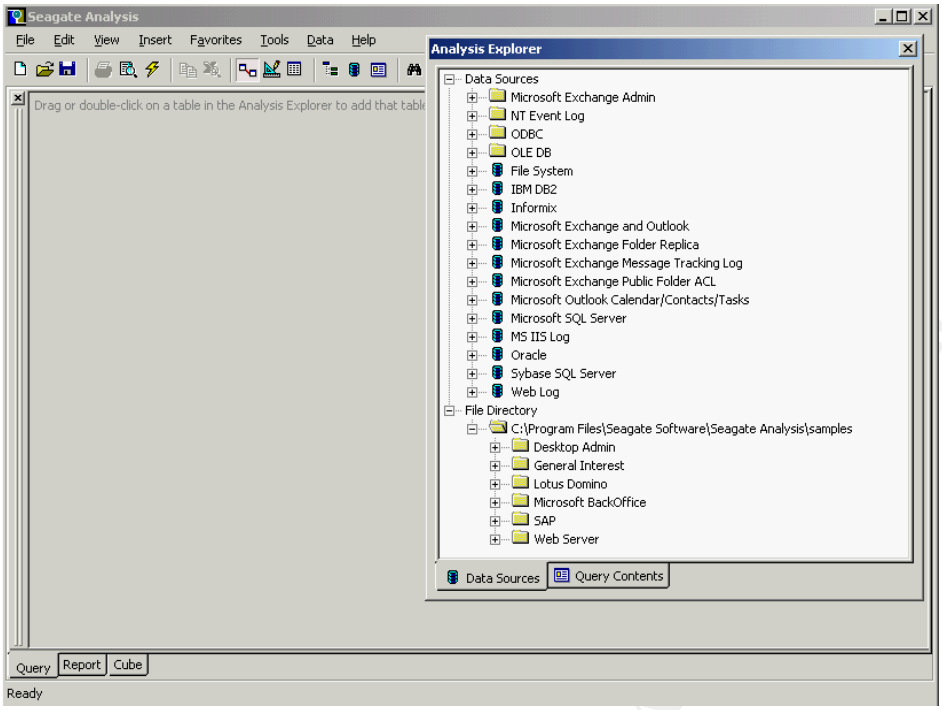

We will start SA on a new analysis document without the Query Assistant. When SA asks for a data source, there is a variety of choices to choose from and we will emphasize on the following data sources:

NT Event Log
This data source allows SA to feed off either archive event logs (.EVT files) or the event logs that are currently kept. The ensuing example will discuss more on how to use this data source.

ODBC (Open Database Connectivity)
ODBC is an application interface standard that allows programmers to abstract their programs from different databases such as MySQL, Microsoft Access and SQL servers (though SA also sports specific drivers to MS SQL and Oracle which are faster than their ODBC brethrens). Each database will have its own ODBC driver that can be installed on the client machine to enable querying from the client machine against the database. More information on ODBC can be obtained from http://www.roth.net/perl/odbc/faq/. For the Windows platform, you would be able to access the ODBC settings through the Control Panel. There is more information on ODBC settings in the reference section.

One word on connecting to data sources: ODBC does a wonderful work in abstracting away the specifics for different databases but there is likely to be kinks that you would have to spend sometime wondering what happened. When things do not connect as smoothly as you wished, try hitting the Crystal Decisions Support page or Google. Chances are high that someone else has already been acquainted with the exact same problem and the issue is solved.

Choosing NT Event Logs will show further choices of the various event logs, which we will choose the Security logs:

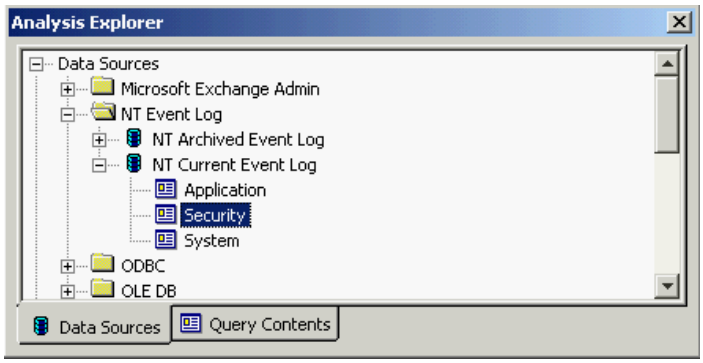

The Security log presents many fields for us to choose and perform analysis on. SA's interface is easy and intuitive to use in this aspect and allows easy selection of the fields that you are interested in. For this example, we choose the following fields:

- Generated Date
- Written Time
- User Name
- Short Description
- Object Name
- Event Type

We will attempt to establish an appreciation of the access patterns that has occurred in this machine through SA. The next screenshot shows the configuration of this analysis document.

SA starts by default in the Query tab. This panel allows the user to create queries against the data source through a point and click interface. The bottom panel shows the data returned from the data source. These returns can be sorted and filtered easily as well. The fields selected from this tab will be used subsequently in building a cube for OLAP analysis.

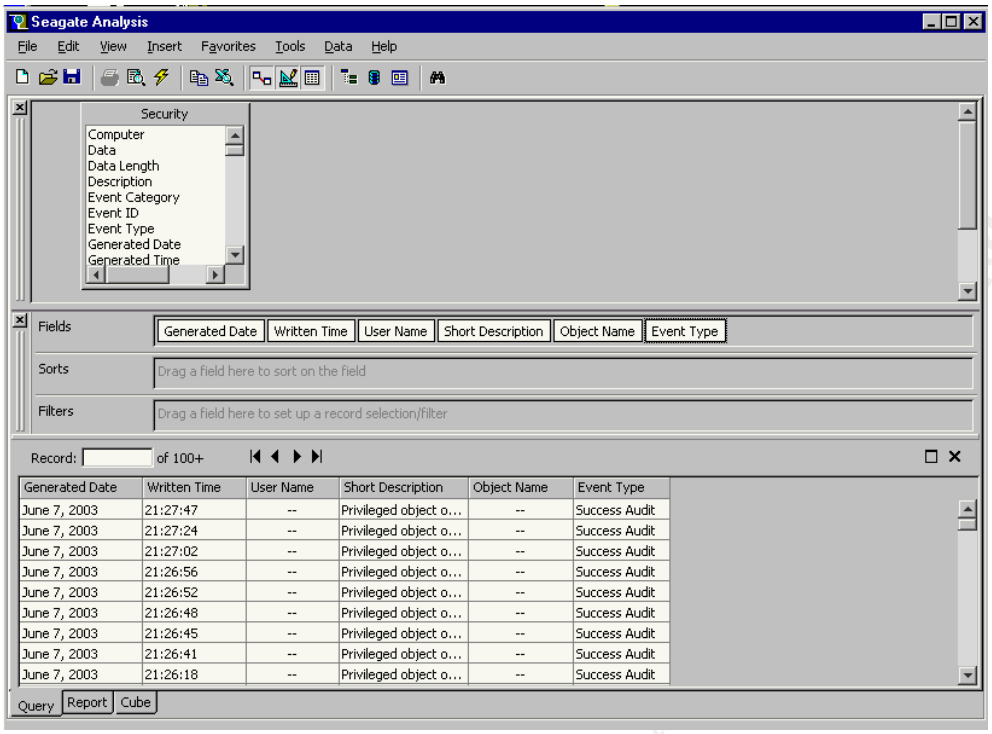

We can proceed to the Report tab by clicking at it on the bottom left of the window. This will bring up the Report panel which we have a screenshot following. The reporting panel sports some rudimentary reporting features and templates. The analyst can use this functionality to produce handsome looking reports. The reports would look similar to many reporting engines that are built into many commercial scanners and log analysis tools as Crystal Decisions' Report engines are commonly used.

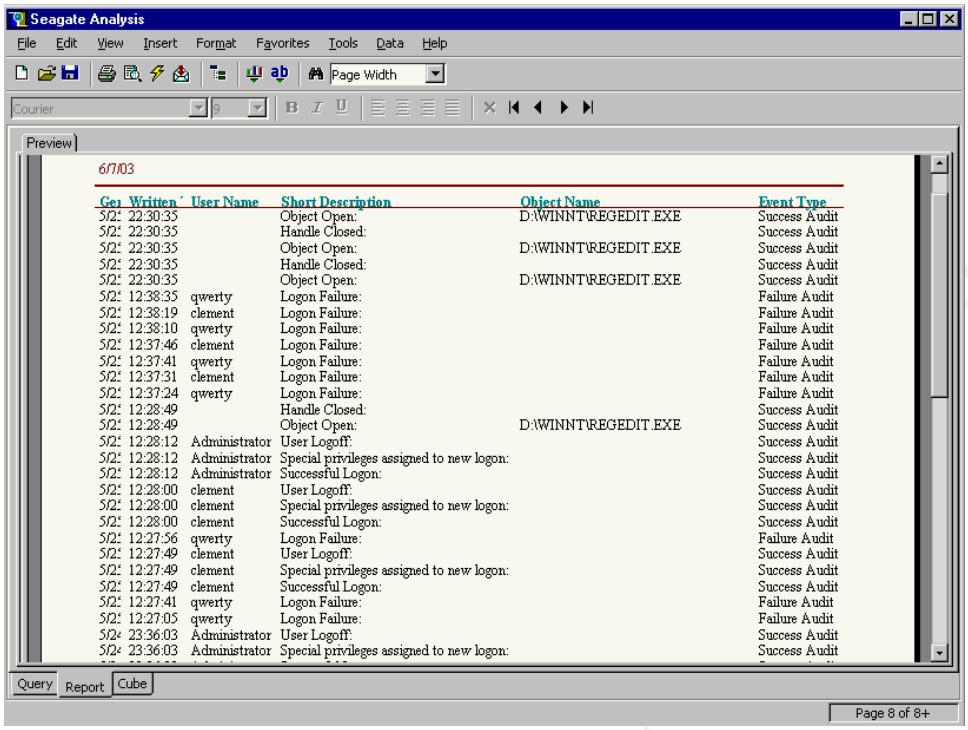

The Cube tab contains the functionality that I consider to be most interesting. This panel allows the analyst to manipulate the cube for viewing the data from various views.

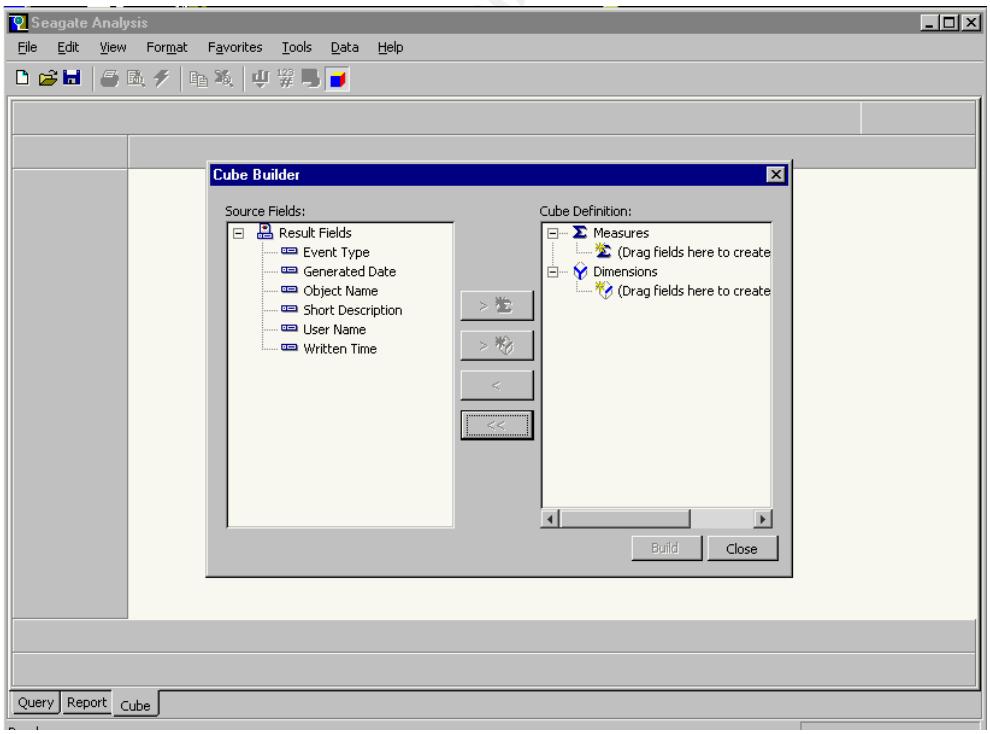

The Cube Builder wizard will appear to guide the user on building a cube. The wizard

prompts the user to identify the fields specified in the Query tab as either the dimensions or measures. The next screenshot shows the various dimensions and measure.

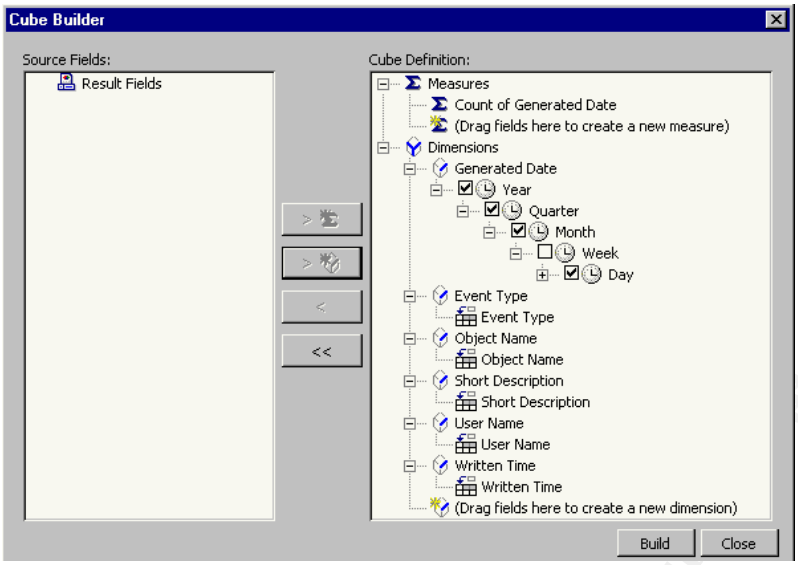

The configuration of the wizard would require some explanations. Since dimensions are objects of an analysis and there are questions to be asked about all the selected fields, all of them thus qualify as dimensions. However, there is a need to generate a proxy to the measure as none of the fields readily qualify as one (i.e. not something that can be an obvious measure such as sales figures in business or number of a particular virus observed). SA is able to automatically count occurrences of a particular field or aggregate it if it is numeric. In this case, we chose 'Generated Date' and since it is not numeric, it becomes a proxy to the number of times a log entry of a particular configuration occurred.

For example, a log entry may be consisted of the following tuple:

{7 June 2003, 21:22:03, Administrator, User Logoff, --, Audit Success}

Since there is no numeric data in this log entry, we can pick a field of the entry and make that a measure. In this case, the 'Generated Date' field is picked. Note that the same field can be still made a dimension even after it is reckoned to be a measure.

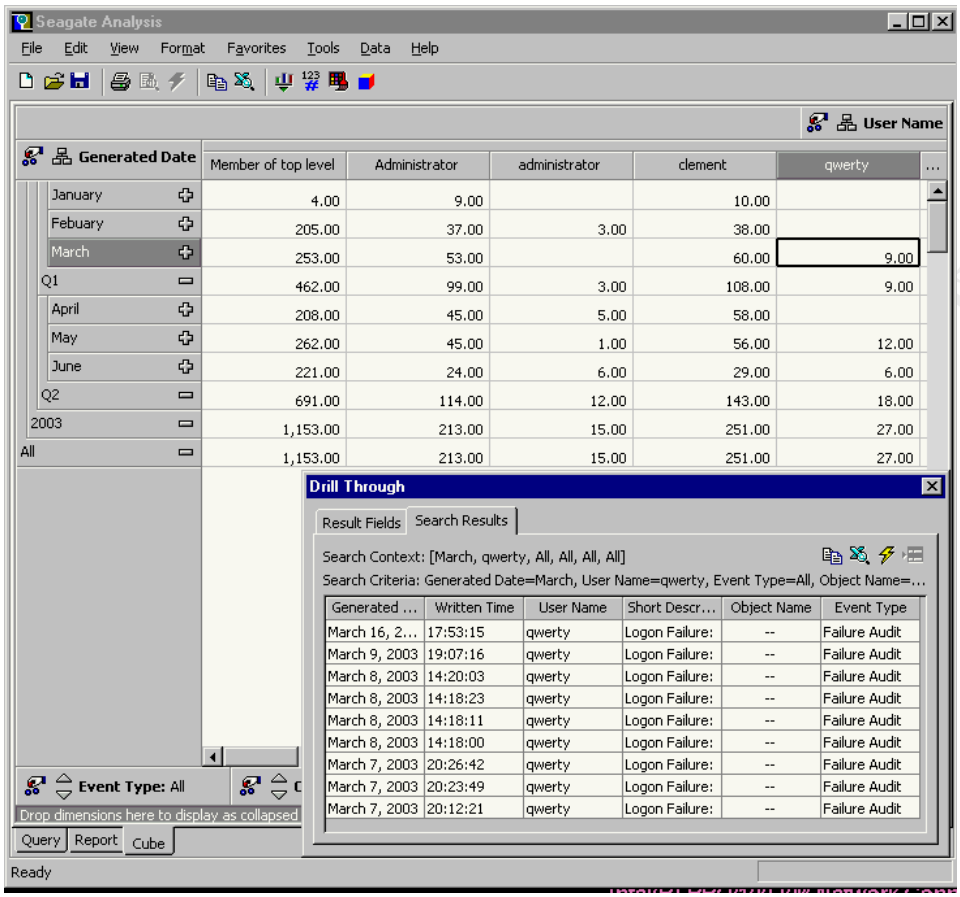

The preceding screenshot appears after the 'Build' button is clicked and SA goes to the data source to gather the data and performs the necessary calculations. A worksheet will then appear with the dimensions lined up at the sides with the corresponding measures. The interpretation of the worksheet will need to be contextualized by the circumstances that the logs were generated. In this case, the measures are counts of NT security events triggered by the respective user accounts grouped by months. Notice that the 'Generated Date' dimension (on the left) is hierarchical and can be expanded or collapsed to reveal or hide the next level of details.

The screenshot also included the 'Drill Through' feature of SA and OLAP tools in general. The selected cell featured in the screenshot is a conjunction of the user 'qwerty' and the month March. The 'Drill Through' option (available through a right click on the cell concerned) shows all the returns pertinent to the chosen cell. This feature allows for deeper examination of trouble spots that the analyst may have inklings about.

We will now discuss more on the slice & dice feature of OLAP. The cube that we have generated is technically called a hypercube since it has more than 3 dimensions. We can see that apart from the lined up dimensions ('Generated Date' and 'User Name'), there are other dimensions are parked at the bottom. These dimensions are interactive and we can slice through each one of them to get to the desired perspective. For example, when

we expand the 'Short Description' dimension, it shows us a more verbose description of the security events that are logged. Rather than displaying the measure for all of these events, we can choose a particular one such as 'Successful Login' and the measures will be recalculated for just that event. Any value in each dimension can be 'sliced' to obtain the desired view of the cube.

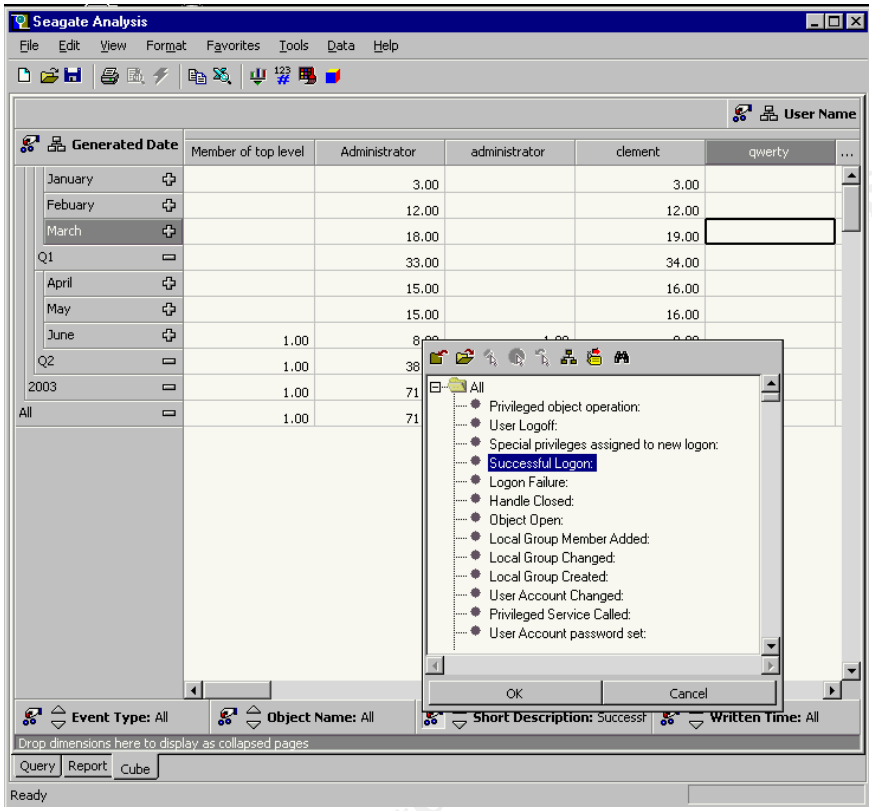

The next interesting feature of SA is its capability to easily allow access to different views of the cube through a simple drag and drop action. The next screenshot will show that by dragging one of the parked dimensions (e.g. 'Short Description' to replace 'User Name') at the bottom of the window to one of the dimensions that are lined up, the analyst will be able to access to another view of the cube quickly. In this case, with an action of drag and drop, the analyst is able to assess the type of security events that are triggered and logged over the months.

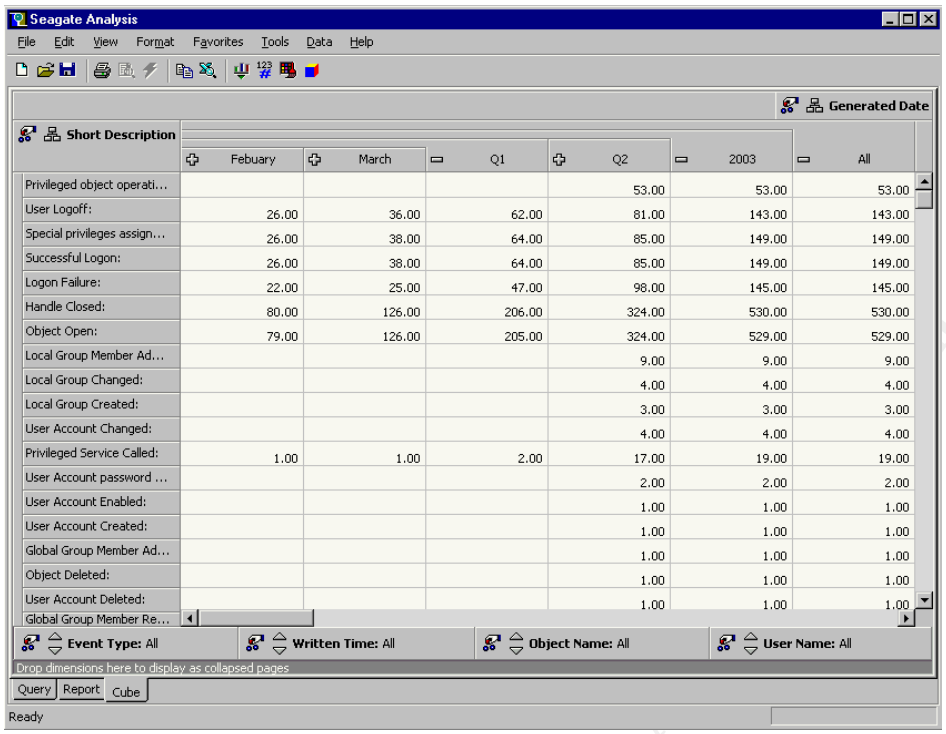

SA also allows the notion of stacked dimension. This means that one dimension is grouped within another and the measures will be calculated and presented in a form similarly grouped. In the following example, we stacked the dimensions 'Short Description' on 'Written Time' and on 'Event Type'. The resultant measures calculated within each grouping tell us the occurrences of successful audits in the area of 'user logoff' broken down by the time of the day. As more dimensions are lined up against the measures, you can see that lesser dimensions are left parked at the bottom for slicing.

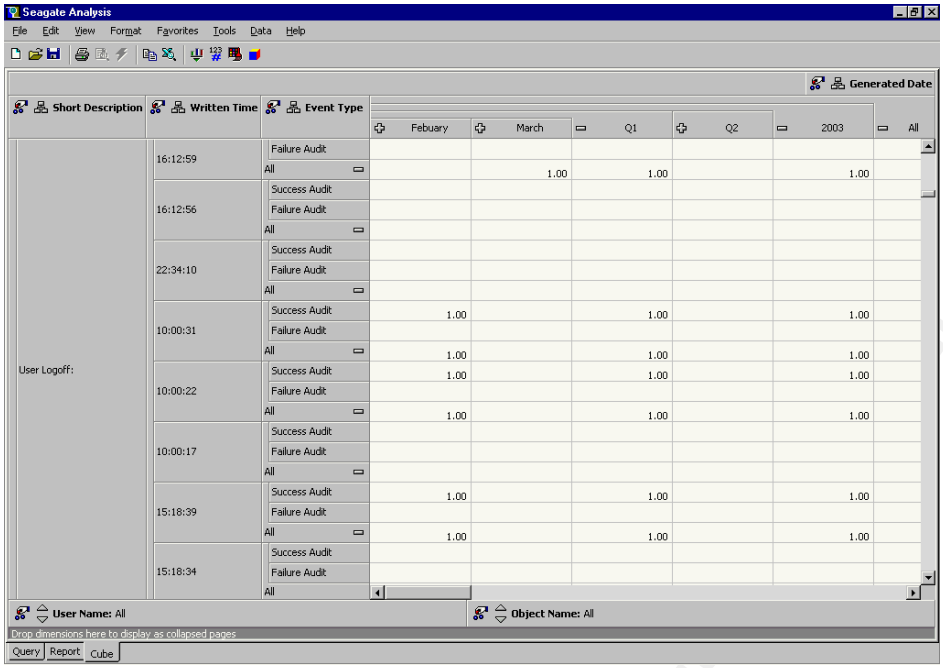

The next functionality that we will discuss is the graphing capability of SA. The graphing capabilities are varied and are constrained to the cells in the worksheet that are graphed. This means that the graph will actively reflect the value that the analyst chooses to graph, similar to Excel's graphing capability. This feature allows interactive examination of the data through graphs. The following screenshot shows a graph depicting the measures mentioned in preceding paragraphs. Graphs are much better in indicating trends and would be a useful tool for the analyst. In particular, SA allows 3D graphs that enable 3 dimensions to be visualized for examination. We will leave the reader to read the cited documentations on SA to discover the wonderful suite of graphing capabilities in SA.

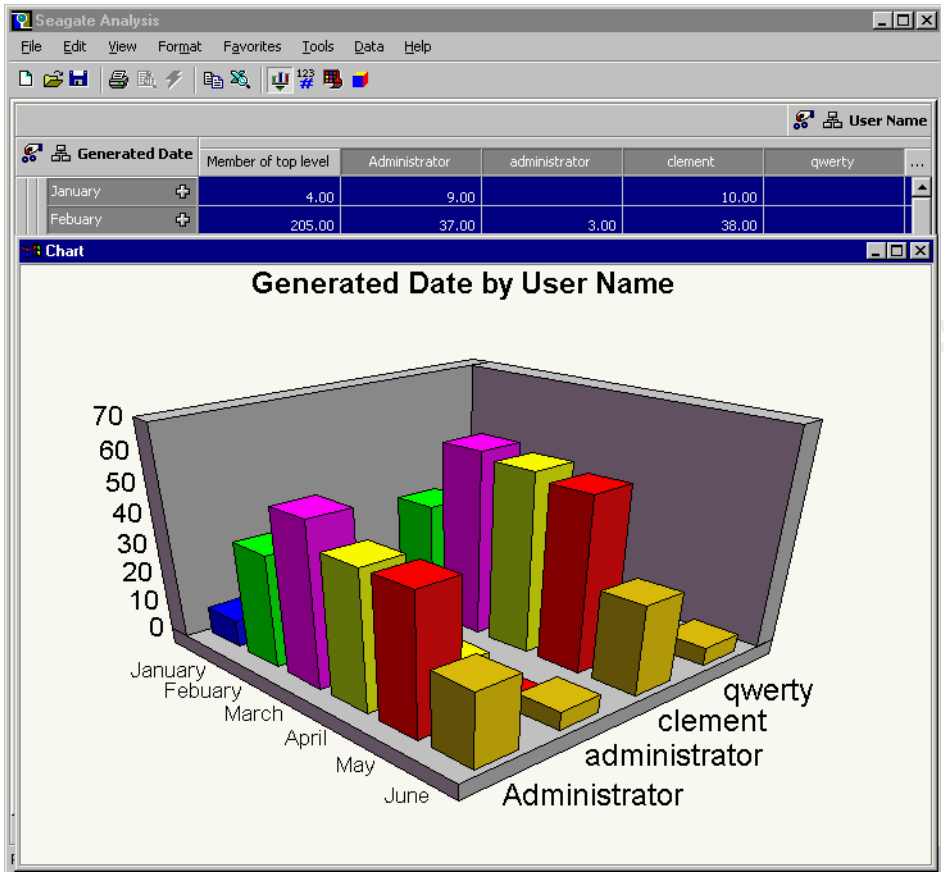

These functionalities are the essence of OLAP's value that SA brings to log analysis. SA does a lot of neat tricks in terms of presentation and nice user interface that makes the process a lot easier and even fun in some instances. In this area, SA outshines other low-end OLAP clients such as the pivot table service provided by Microsoft Excel. The reader is encouraged to explore the various documentations cited to better appreciate SA's offerings. One other feature that is worth mentioning is the capability to readily export a worksheet to Excel format. This would then allow an analyst familiar with Excel to perform subsequent analysis on Excel if needed.

Database Schema Considerations

The excursion through the NT Event Logs via SA showed only half the story in an OLAP approach to log analysis. In many cases where SA can be useful, it would be used against a database where its full power can be harnessed. Years of OLAP research has identified that the database should adopt a schema pattern, as it is most friendly to OLAP operations without unnecessarily compromising the speed and efficiency of database queries.

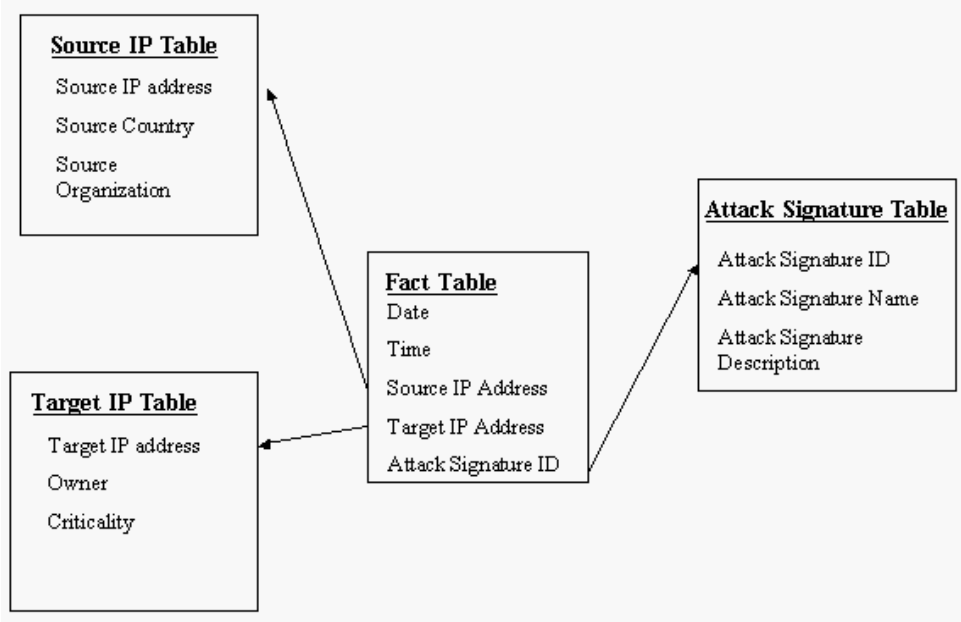

The above diagram shows a possible star schema for the 2 cubes that we used as examples to open the discussion on OLAP. The star schema is named as such because the resulting schema diagram shows a central table radiating outwards to other tables, resulting in a star-like look. If you are not familiar with relational database, there are some references cited at the end of this paper that will be useful. For the more SQL savvy, a fact table would essentially be holding foreign keys to the various dimension tables such as the 'Source IP', 'Target IP' and 'Attack Signature' tables.

The fact table hence holds the data entered into this database. Whenever a new event has occurred, logged and entered into the database, it will be entered into the fact table. As the contents of the fact table are concise, the dimension tables will be used to flesh out further details of the events. There are a lot more other considerations in designing an optimal schema for OLAP processing, do look at the references cited at the end for further information.

Particularly for SA, when an analysis document is using a database that sports a star schema, the user is able to choose the relevant tables to be included as data sources. The

user then either link up the various keys using the 'Smart Link' feature that automatically links up the keys or perform the linking manually through drag and drop. The user is then able to perform analysis on a richer set of fields available through the dimension tables. For example in this schema, we can obtain distribution of attackers in terms of originating countries via the 'Source Country' field in the 'Source IP' table. We can also find out how many high criticality servers are attacked by performing analysis on the 'Criticality' field in 'Target IP' table.

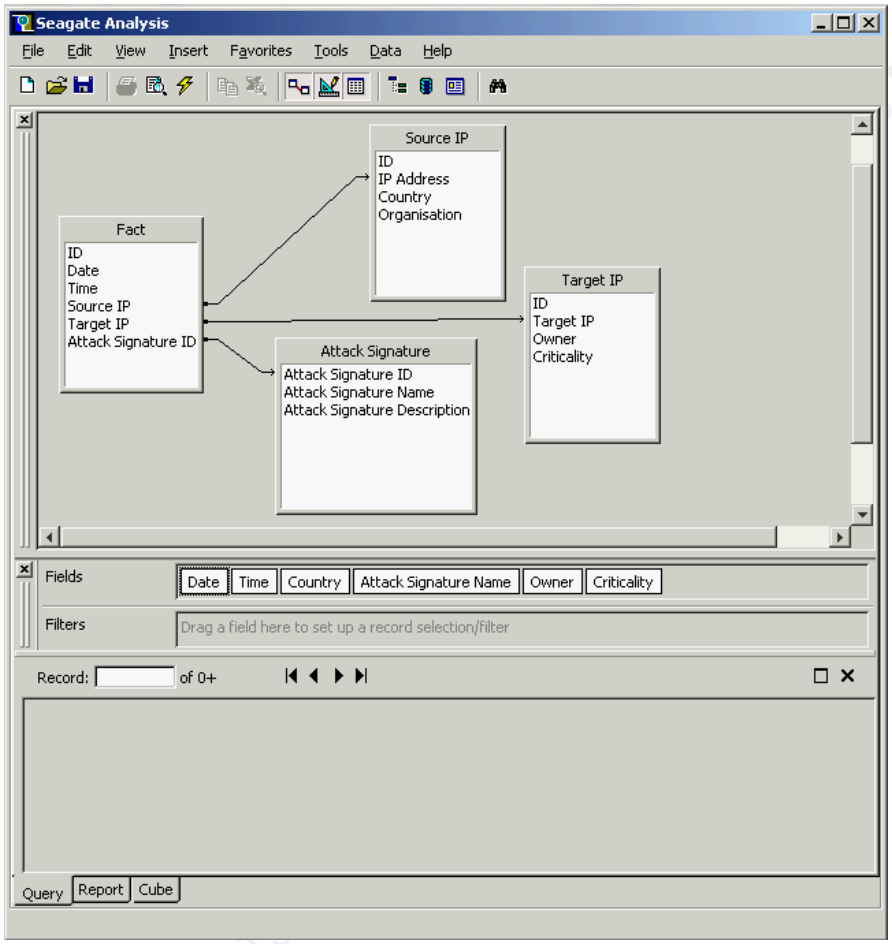

The above screenshot shows how the star schema can be used in SA. The above analysis document is built as described using a Microsoft Access database file as data source. The tables were chosen and the link relationships are defined manually through a drag-and-drop interface. Notice the fields that are chosen for display. Though the fact table is the one continuously fed with new logs, not all of its fields will be displayed. Some of the fact table's fields are linked to particular records in the dimension tables and the chosen corresponding fields will be displayed instead. So in this case, where the displayed fields are:

{Date, Time, Country, Attack Signature Name, Owner, Criticality}

There is no report of the IP addresses at all and an analysis base on larger entities (such as source country rather than source IP address) can be performed instead. In fact, the star schema is a special case of a more general schema pattern called the snowflake schema which can be used for more complicated data model.

An Analyst's Swiss Army Knife

Just as *netcat* is featured as a Swiss army knife for the hackers and penetration testers, OLAP clients such as SA emerge as a flexible tool for the analyst. Dedicated OLAP clients allow users that are well versed in OLAP concepts to investigate large sets of logs in ways that are probably not sufficiently addressed by any other tools. Log analysis tools also attempt to abstract processing details (which could be quite akin to OLAP) and often result in canned metrics. Though there are approaches such as using Perl scripts to sieve out suspicious log entries via regular expressions, I believe that along with other benefits of managing logs via databases[1], analysis of logs via dedicated OLAP clients represents a lightweight and yet powerful means of analysis.

There were many allusions to the utility of OLAP in log analysis for security purposes peppered throughout this paper, and many may feel in their gut that there is a significant role to be filled by OLAP and tools like SA. Can we list some examples to convince the skeptics among us? This is notwithstanding that commercial log analysis tools are already bundling OLAP functionalities and hopefully not just for marketing value arising from the cool 3D graphs:

- Monitoring of thresholds is a common tactic to police a policy. For example 10 failed logouts in a month may be flagged as suspicious. SA has a feature under the Cube panel called Highlighting which allows specification of a pair of high and low values to indicate the threshold. Upon specification of this pair of values, the cells in the worksheet will light up in green, amber of red to identify each measures standing within the threshold.
- Graphing of measures over time can be telling on whether there are seasonality factors at play. For example, if the object of analysis is the consolidated logs from all personal firewalls deployed in an enterprise, it would be possible to assess how many unfamiliar programs are attempting outbound connections, whether they are all connecting to similar locations and whether there are more instances of such connections in a particular month, say November maybe because the program is activated on Halloween's Day. I would be hard pressed to figure that out without a graph.
- OLAP's crown jewel of multidimensional viewing to aid analysis can be very

[1] While it is not a trivial undertaking to convert existing textual logs into databases, the advent of using database as data store for logging platforms such as SYSLOG (Syslog-ng, mysyslog, Kiwi, WinSyslog etc.) should alleviate this problem. Interestingly, Microsoft released a tool called LogParser that allows SQL like querying against textual log files instead, offering an alternative to querying such log files using Perl. Among many other things that LogParser does, one feature that is of particular pertinence here is the capability to take a large variety of textual log files and output them as relational tables in a database.

useful in tackling situations with several parameters. For example, assume access logs were kept for a web authentication page and the user name, IP address, browser type, time and login status were recorded. We can create a hypercube out of these dimensions and see whether there are uncharacteristically high counts of login failures for a particular user from a particular IP address at a certain period of the day using a particularly curious brand of browser. The presence of such traits may indicate a particularly insistent attacker, slowly trying out a few passwords at a time. Incidentally, data mining is the more automated way to achieve this endeavor.

- Just for the ease of access to log analysis. OLAP tools have evolved to a stage that it is meant to be use by non-technical people. This is because the tool is meant to serve reporting and some analysis needs for sales persons, line managers and senior managers in the business world. Using OLAP tools on log analysis for security will lower the barrier for practitioners to really peer at their logs.

In spite of these possibilities, it still remains that the value of OLAP tools in security log analysis depends heavily on the strategy adopted in collecting the right logs at the right granularity for the right kind of analysis. However, armed with your thinking cap and a bunch of logs that you've collected over the months, OLAP tools such as SA should be able to give you new insights in a jiffy without much planning as well.

Conclusion

This paper introduces the basic tenets of OLAP, particularly in their application to the log analysis problem. The methods described in this paper are relatively unsophisticated but they are chosen to illustrate the main concepts. The availability of such powerful tools should already fire up much curiosity of what information could be hidden in their stash of logs in security professionals who do not want pricey log analysis products.

One point to be made on the robustness of SA: we have used it to look at logs amounting to 2 million records and given the necessary processing time, SA worked fine. The ad hoc mode of OLAP that SA is used as it is described in this paper can take quite a bit of time if large amount of data is spewed out by the database and stream to SA over the network for further processing. The field of OLAP has developed an entire branch of technologies to ease such pain for the impatient.

There is no reason why SA should be the security practitioner's choice of OLAP client. The reference cites a survey of the various OLAP clients available and how they measure up. However, SA's availability in terms of licensing and features are unmatched at its price range. Through this paper, I hoped to illustrate what value OLAP can bring to security log analysis and attempted to jump start the fellow practitioners' understanding of OLAP concepts in order to be readily productive in using the technology. Nonetheless, other than allowing me to clearly illustrate the value of OLAP in this paper, SA's status as a free tool has ramifications on how readily can OLAP be used by security

practitioners, though significantly limited by its rarity on the Web.

What is needed is an OLAP client that is as free as other well-known security tools such as SNORT, SYSLOGD, IPCHAINS or even MySQL. The availability of such tools has long encouraged spontaneous experimentations of new ways to cost-effectively tackle the security problem. If novelty proved to be elusive, such tools at least offer the freedom to partner their respective efficacies in ways unfettered by proprietary concerns and only limited by one's imagination of what's possible. However, for SA's case, Crystal Decisions has stopped the SA-for-free program in 2002. Though if you do have the Dog CD, which the free SA distribution had been affectionately called, stashed somewhere among your CD wallets, it would be worthwhile checking it out in this new light.

A word on OLAP technologies in the open source scene: in my search so far it seems that development of OLAP tools is not vibrant under the open source scheme. Many tools are either software components such as ActiveX controls or JSP tag library. Mondrian (http://apoptosis.dyndns.org:8080/open/mondrian/doc/index.html) and its client JPivot figured quite prominently in this field as open source OLAP tools but they do not really come as ready-to-use a package as SA does.

More security professionals are going to pump logs into databases, and OLAP will be the natural technology to embrace to get more mileage from their logs. The Seagate Analysis tool represents a great opportunity for security professionals to explore, in an unfettered manner, the wide field of OLAP and in the process empowering themselves to be better informed by their logs.

Who knows, they may yet discover that trees are indeed falling in their forests.[2]

References

OLAP	OLAPReport – What is OLAP? http://www.olapreport.com/fasmi.htm OLAP Council White Paper http://www.olapcouncil.org/research/whtpaply.htm OLAP Clients for Microsoft OLAP Services http://www.sqlpass.org/voting/sessions/99compass/compass_html/109/sld001.htm

[2] The subtitle 'A Cube to Rule Them All' is an allusion to that of 'The Lord of the Rings – A Ring to Rule Them All'. This initial muse struck me more profoundly in the concluding statement. The novel spoke of a dying species called *Ents* which are tree-herders: they too in a way try to prevent trees falling in their forests.

SQL and Relational Database Design	A Gentle Introduction to SQL http://sqlzoo.net/ Giovinazzo, A. Williams. Object-oriented Data Warehouse Design: Building a Star Schema, Prentice Hall, 2000 Introduction to Relational Database Design http://www.edm2.com/0612/msql7.html Mark Levene, George Loizou. Object-oriented Data Warehouse Design: Building a Star Schema, 1999
Log Analysis	Counterpane: Log Analysis Resources http://www.counterpane.com/log-analysis.html Shmoo LogAnalysis mailing list http://sisyphus.iocaine.com/pipermail/loganalysis/ Advanced Log Processing, Anton Chuvakin http://www.securityfocus.com/infocus/1613
Database Connectivity	Setting Up Your Database Connection Using ODBC http://support.microsoft.com/default.aspx?scid=/support/frontpage/fp2000/aspweb/page00003.asp
Log and Database	Logging Syslog to a Database, Zbyszek Sobiecki http://ezine.daemonnews.org/200111/syslog.html Syslog-ng http://www.balabit.com/products/syslog_ng/ WinSysLog http://www.winsyslog.com/en/FAQ/Setup-WinSyslog-with-MySql.asp Microsoft Log Parser version 2.0 http://www.microsoft.com/windows2000/downloads/tools/logparser/default.asp

Example Log Analysis Products	Intellitactics' Network Security Manager http://www.infosecuritymag.com/2002/oct/testcenter.shtml A Technically Better Network Intrusion Detection System, Intrusion Inc. https://www.intrusion.com/products/downloads/Intrusion_SecureNet_Overview_12202002.pdf ZoneLog Analyzer http://zonelog.co.uk/ Sawmill Log Analyzer http://www.sawmill.net
Seagate Analysis and Seagate Info	What is Seagate Analysis? http://support.crystaldecisions.com/library/docfiles/sa/en/ Crystal Decisions Forum on Seagate Analysis http://support.crystaldecisions.com/forums/content.asp?fid=5&sk=5&ps=100& Crystal Decisions Support at http://support.crystaldecisions.com/library/kbase.asp?ref=default.asp_selectlist