Real-time log file analysis using the Simple Event Correlator (SEC)

John P. Rouillard
University of Massachusetts at Boston
For LISA 2004 Conference: Atlanta, GA

November 2004

Abstract

Log analysis is an important way to keep track of computers and networks. The use of automated analysis always results in false reports, however these can be minimized by proper specification of recognition criteria. Current analysis approaches fail to provide sufficient support for the recognizing the temporal component of log analysis. Temporal recognition of event sequences fall into distinct patterns that can be used to reduce false alerts and improve the efficiency of response to problems. This paper discusses these patterns while describing the rationale behind and implementation of a ruleset created at the CS department of the University of Massachusetts at Boston for SEC - the Simple Event Correlation program.

1 Introduction

With today's restricted IT budgets, we are all trying to do more with less. One of the more time consuming, and therefore neglected, tasks is the monitoring of log files for problems. Identifying and resolving these problems is the first step in keeping one's systems, users, and bosses happy and healthy. Failure to identify and resolve these problems quickly leads to downtime and loss of productivity. Log files can be verbose with errors hidden among the various status events indicating normal operation. For a human, finding errors among the routine events can be difficult, time consuming, boring, and very prone to error. This is exacerbated when aggregating events, using a mechanism such as syslog, due to the intermingling of events from different hosts that can submerge patterns in the event streams.

Many monitoring solutions rely on summarizing the log files for the previous days logs. This is very useful for accounting and statistics gathering. Sadly, if the goal is problem determination and resolution then reviewing these events the day after they are generated is less helpful. Systems administrators cannot proactively resolve or quickly respond to problems unless they are aware that there is a problem. It is not useful to find out in the next morning's summary that a primary NFS server was reporting problems five minutes before it went down. This is especially bad if the critical server takes a dozen other systems down with it. The sysadmin staff needs to discover these problems while there is still time to fix the problem and avert a catastrophic loss of service.

The operation of computers and computer networks evolves over time and requires a solution to log file analysis that address this temporal nature. This paper describes some of the current issues in log analysis and presents the rationale behind an analysis rule set developed at the Computer Science Department at the University of Massachusetts at Boston. This ruleset is implemented for the Simple Event Correlator (SEC), which is a Perl based tool designed to perform analysis of plain text logs.

1.1 Current Approaches

There are many programs that try to isolate error events by automatically condensing or eliminating

routine log entries. In this paper, I do not consider interactive analysis tools like MieLog[Takada02]. I separate automatic analysis tools into offline or batch monitoring and on-line or real-time monitoring.

1.1.1 Offline Monitoring

Offline solutions include: **logwatch**[logwatch], **SLAPS-2**[SLAPS-2], or **Addamark LMS**[Sah02]. Batch solutions have to be invoked on a regular basis to analyze logs. They can be run once a day, or many times an hour. Offline tools are useful for isolating events for further analysis by real time reporting tools. In addition they provide statistics that allow the system administrator to identify the highest event sources for remedial action. However, offline tools do not provide the ability to provide automatic reactions to problems. Adam Sah in discussing the **Addamark LMS**[Sah02, p. 130] claims that real-time analysis is not required because a human being, with slow reaction times, is involved in solving the problem. I disagree with this claim. While it is true that initially a human is required to identify, isolate and solve the problem, once it has been identified, it is a candidate for being automatically addressed or solved. If a human would simply restart apache when a particular sequence of events occur, why not have the computer automatically restart apache instead? Automatic problem responses coupled with administrative practices can provide a longer window before the impact of the problem is felt. An "out of disk space" condition can be addressed by removing buffer files placed in the file system for this purpose. This buys the system administrator a longer response window in which to locate the cause of the disk full condition minimizing the impact to the computing environment.

Most offline tools do not provide explicit support for analyzing the log entries with respect to the time they were received. While they could be extended to try to parse timestamps from the log messages, this is difficult in general, especially with multiple log files and multiple machines, as ordering the events requires normalizing the time for all log messages to the same timezone. Performing analysis on log files that do not have timestamps eliminates the ability of these batch tools to perform analysis by time. Solutions such as the **Addamark LMS**[Sah02] parse and record the generation time, but the lack of real-time event-driven, as opposed to polled, triggers reduces its utility.

1.1.2 Online Monitoring

Online solutions include: **logsurfer**[logsurfer], **logsurfer+**[logsurfer+], **swatch**[swatch, Hansen1993], **2swatch**[2swatch], **SHARP**[Bing00], **ruleCore**[ruleCore], **LoGS**[LoGS] and SEC[SEC]. All of these programs run continuously watching one or more log files, or receiving input from some other program.

Swatch is one of the better known tools. **Swatch** provides support for ignoring duplicate events and for changing rules based on the time of arrival. However, **swatch's** configuration language does not provide the ability to relate arbitrary events in time. Also, it lacks the ability to activate/deactivate rules based on the existence of other events other than suppressing duplicate events using its throttle action.

Logsurfer dynamically changes its rules based on events or time. This provides much of the flexibility needed to relate events. However, I found its syntax difficult to use (similar to the earliest 1.x version of SEC) and I never could get complex correlations across multiple applications to work properly. The dynamic nature of the rules made debugging difficult. I was never able to come up with a clean, understandable, and reliable method of performing counting operations without resorting to external programs. Using SEC, I have been able to perform all of the operations I implemented in **logsurfer** with much less confusion.

LoGS is an analysis program written in Lisp that is still maturing. While other programs create their own configuration language, **LoGS's** rules are also written in Lisp. This provides more flexibility in designing rules than SEC, but may require too much programming experience on the part of the rule designers. I believe this reduces the likelihood of its widespread deployment. However, it is an exciting addition to the tools for log analysis research.

The Simple Event Correlator (SEC) by Risto

Vaarandi uses static rules unlike **logsurfer**, but provides higher level correlation operations such as explicit pair matching and counting operations. These correlation operations respond to a triggering event and persist for some amount of time until they timeout, or the conditions of the correlation are met. SEC also provides a mechanism for aggregating events and modifying rule application based on the responses to prior events. Although it does not have the dynamic rule creation of **logsurfer**, I have been able to easily generate rules in SEC that provide the same functionality as my **logsurfer** rules.

1.1.3 Filter in vs. filter out

Should rules be defined to report just recognized errors, or should routine traffic be eliminated from the logs and the residue reported? There appear to be people who advocate using one strategy over the other.

I claim that both approaches need to be used and in more or less equal parts. I am aware of systems that are monitored for only known problems. This seems risky as it is more likely an unknown problem will sneak up and bite the unwary systems administrator. However, very specific error recognition is needed when using automatic responses to ensure that the best solution is chosen. Why restart apache if killing a stuck cgi program will solve the problem?

Filtering out normal event traffic and reporting the residue allows the system administrator to find signatures of new unexpected problems with the system. Defining "normal traffic" in such a way that we can be sure its routine is tricky especially if the filtering program does not have support for the temporal component of event analysis.

I have seen filters that ignore abnormal reboot traffic as a side effect of trying to ignore the normal reboot traffic. One system was scheduled to reboot at 6PM and had an extensive list of "normal traffic". These events were filtered out of the data stream so that alerts were not raised for a normal scheduled occurrence. The data processing run started at 6:30PM, and the system rebooted unexpectedly around 9PM. This reboot and subsequent loss of the analyzed data was not detected until 7AM the next morning when an expected "analysis done" event was found to be missing. If the filter had provided the ability to to ignore reboot events during just the 6-6:15 PM window, or it reported reboot events that occurred during a data processing run, the operator could have restarted the data analysis getting the data to the department with a less of a delay.

```
>  CMD: /usr/lib/sendmail -q
>  root 25453 c Sun May 23 03:31:00 2004
<  root 25453 c Sun May 23 03:31:00 2004
```

Figure 1: Events indicating normal activity for a scheduled cron job.

2 Event Modeling

Modeling normal or abnormal events requires the ability to fully specify every aspect of the event. This includes recognizing the content of the event as well as its relationship to other events in time. With this ability, we can recognize a composite or correlated event that is synthesized from one or more primitive events. Normal activity is usually defined by these composite events. For example a normal activity may be expressed as:

> '`sendmail -q`' is run once an hour by root at 31 minutes after the hour. It must take less than 1 minute to complete.

This activity is shown by the cron log entries in Figure 1 and requires the following analysis operations:

- Find the sendmail CMD line verifying its arrival time is around 31 minutes after the hour. If the line does not come in, send a warning.
- The next line always indicates the user, process id and start time. Make sure that this line indicates that the command was run by root.
- The time between the CMD line arrival and the final line must be less than 1 minute. Because other events may occur between the start and

end entries for the job, we recognize the last line by its use of the unique number from the second field of the second line.

This simple example shows how a tool can analyze the event log in time. Tools that do not not allow the specification of events in the temporal realm as well as in the textual/content space can suffer from the following problems:

- matching the right event at the wrong time. This could be caused by an inadvertent edit of the cron file, or a clock skew on the source or analyzing host.
- not noticing that the event took too long to run.
- not noticing that the event failed to complete at all.

2.1 Temporal relationships

The cron example mentions one type of temporal restriction that I call a schedule restriction. Schedule restrictions are defined by working on a defined (although potentially complex) schedule. Typical schedule restrictions include: every hour at 31 minutes past the hour, Tuesday morning at 10AM, every weekday morning between 1 and 3AM.

In addition to schedule restrictions, event recognition requires accounting for inter-event timing. The events may be from a single source such as the sequence of events generated by a system reboot. The statement that the first to last event in a boot sequence should complete in five minutes is an inter-event timing restriction. Also, events may arise from multiple sources. Multi-source inter-event timing restrictions might include multiple routers sending an SNMP authentication trap in five minutes, or excessive "connection denied" events spread across multiple hosts and multiple ports indicating a port scan of the network.

These temporal relationships can be explicit within a correlation rule: specifying a time window for counting the number of events, suppressing events for a specified time after an initial event. The timing relationship can also be implicit when one event triggers the search for subsequent events.

2.2 Event threading

Analysis of a single event often fails to provide a complete picture of the incident. In the example above, reporting only the final cron event is not as useful as reporting all three events when trying to diagnose a cause. Lack of proper grouping can lead to underestimating the severity of the events. Consider the following scenario:

1. A user logs in using ssh from a location that s/he has never logged in from before.
2. The ssh login was done using public key authentication.
3. The ssh session tries to open port 512 on the server. It is denied.
4. Somebody tries to run a program called "crackme" that tries to execute code on the stack.
5. The user logs out.

Looking at this sequence implies that somebody broke in and tried to execute an unsuccessful attempt to gain root privileges. However, in looking at individual events, it is easy to miss the connections.

- A user logs in using ssh from a location that s/he has never logged in from before.

  Login/logouts are routine traffic that we want to filter out unless something strange happens. Is the login from a new location a problem by itself? This depends on local policy, but probably not. However, this information is needed to determine who logged in, in case something weird does happen.

- The ssh login was done using public key authentication.

  Useful info if the user never used public key before and always used password authentication in the past. However, it may easily be dismissed as the person may be learning to use public key authentication. This event again falls under routine traffic.

- The ssh session tries to open port 512 on the server. It is denied.

  This is a bit weird, but did somebody just commit a typo and mean 5912, a VNC server port, or did they really want to bind to the exec port?.

  This ssh event does not include identifying information, specifically it does not include the user or the source address for the ssh process. The sshd event is labeled with the PID of the child sshd process, so the administrator can try to manually correlate the child PID to the parent process's login event that reports the username.

- Somebody tries to run a program called "crackme" that tries to execute code on the stack.

  We do obtain a username from here, but is it a (poorly named) CS 240 project that ran amok on a bad pointer dereference, or is it a hacking attempt?

- The user logs out.

  Did the user immediately log out after running the program in an attempt to escape detection? Did the user stay logged in for a while to try to debug the program? Logouts again fall into the routine traffic category.

Taken in isolation, each event could be easily dismissed, or even filtered out of the reports. Reporting them as discrete events, as many analysis tools do, may even contribute to an increased chance of missing the pattern. Taken together they indicate a problem that needs to be investigated. A log analysis tool needs to provide some way to link these disparate messages from different programs into a single thread that paints a picture of the complete incident.

2.3 Missing Events

Log analysis programs must be able to detect missing log events[Finke2002]. These missing events are critical errors since they indicate a departure from normal operation that can result in a many problems. For example, cron reports the daily log rotation at 12:01AM. If this job is not done (say, because cron crashed), it is better to notice the failure immediately rather than three months later when the partition with the log files fills up.

The problem with detecting missing events is that log monitoring is – by its nature – an event-driven operation: if there is no event, there is no operation. The log analysis tool should provide some mechanism for detecting a missing event. One of the simpler ways to handle this problem is to generate an event or action on a regular basis to look for a missing event. An event-driven mechanism can be created using external tools such as cron to synthesize events, but I fail to see a mechanism that the log analysis tool can use to detect the failure of the external tool to generate these events.

2.4 Handling False Positives/False Negatives

A false negative occurs when an event that indicates a problem is not reported. A false positive results when a benign event is reported as a problem. False negatives impact the computing environment by failing to detect a problem. False positives must be investigated and impact the person(s) maintaining the computing environment. A false positive also has another danger. It can lead to the "boy who cried wolf" syndrome, causing a true positive to be ignored as a false positive.

Two scenarios for generating false negatives are mentioned above. Both are caused by incorrectly specifying the conditions under which the events are considered routine.

False positives are another problem caused by insufficiently specifying the conditions under which the event is a problem. In either case, it may not be possible to fully specify the problem conditions because:

- Not all of the conditions are known.

- Some conditions are not able to be monitored and cannot be added to the model.

It may be possible to find correlative conditions that occur to provide a higher degree of discrimination in the model. These correlative events can be

use to change the application of the rules that cause the false positive to inhibit the false report.

I had a problem many years ago where one router in a redundant pair of routers would effectively loose its routing entries. Monitoring was in place to report routing changes on the routers. However, no routing changes were reported. We were able to detect the problem by watching the throughput of the downstream interfaces. However, variations in traffic, and background routing updates and other traffic interfered with the ability to detect the problem using this method. As a correlative event, we set up a sniffer to watch for router redirect events. When the router redirect events were detected without a corresponding routing change event, we rebooted the router.

To reduce these false positives and false negatives, the analysis program needs to have some way of generating and receiving these correlative events.

Above, I claimed that it is currently impossible to eliminate all false events. I have generated false events during my manual log analysis. Even though computers do appear more intelligent than I am, I contend that this is an illusion. Computers will have to exhibit far greater intelligence than humans to eliminate false events. However, by proper specification of event parameters, false events can be greatly reduced.

2.5 Single vs. Multiple Line Events

Programs can spread their error reports across multiple lines in a logfile. Recognizing a problem in these circumstances requires the ability to scan not just a single line, but a series of lines as a single instance. The series of lines can be treated as individual events, but key pieces of information needed to trigger a response or recognize an event sequence may occur on multiple lines. Consider the cron example of Figure 1: the first two lines provide the information needed to determine that it is an entry for sendmail started by root, and the process id is used in discovering the matching end event. Handling this multi-line event as multiple single line events complicates the rules for recognizing the events.

Multi-line error messages seem to be more prevalent in application and device logs that do not use the Unix standard syslog reporting method, but some syslog versions split long syslog messages into multiple parts when they store them in the logfile. Fortunately, when I have seen this happen, the log lines always occur adjacent to one other without any intervening events from other sources. This allows recognition provided that the split does not occur in the middle of a field of interest.

With syslog and other log aggregation tools, a single multi-line message can be distorted by the injection of messages from other sources. The logs from applications that produce multi-line messages should be directed to their own log file so that they are not distorted. Then a separate SEC process can analyze the log file and create single line events that are passed to a parent SEC for global correlation. This is similar to the method used by **Addamark**[Sah02].

Although keeping the log streams separate simplifies some log analysis tasks, it prevents the recognition of conditions that affect multiple event streams. Although SEC provides a mechanism for identifying the source of an event, performing event recognition across streams requires that the event streams be merged.

2.6 State persistence

Since log monitoring and correlation is the result of tracing the various states of the system, we would like some way to initialize the state information quickly so that we can begin our correlation operations as soon as possible after startup or reconfiguration of the log analysis tool.

Loading a saved state upon startup is dangerous unless we have a mechanism that can verify that the state we loaded matches the current state of the system. The longer the period of time between the state being saved and being reused, the greater the chance that we no longer have a valid state and will make mistakes using these facts to make decisions.

Systems such as HPOV NNM have a "settling time" where they learn the state of the network. This learning interval can be found in many tools and I believe it increases with the complexity of the correlation operation, since more complex correlations require more state information that takes longer to

gather. Correlation rules are influenced by the context or state information that is available to them. The correlations then change their operation based on their context. When the correlation engine starts up, rules to detect the occurrence of three events in a 5 minute window may not fire, because there is no longer a contextual record of one of the events that was processed before the engine shutdown. Hence the correlation engine has made an incorrect decisions because it has not yet learned enough context to make the correct decision.

A similar problem occurs with active correlation operations. On correlation engine restarts, active correlation operations are no longer present. For example, Event1 is processed. This starts a correlation operation that takes some action and suppresses Event1 for the next two hours. Ten minutes later, the correlation engine is restarted. After the software restarts, the active correlation operation is not active anymore, so the next occurrence of Event1 takes some actions and activates the suppression rule even though it should not have had any effect.

SEC has the ability to perform a soft reconfiguration that preserves the state of contexts while the state of active correlations are lost.

3 Sec correlation idioms and strategies

This section describes particular event scenarios that I have seen in my analysis of logs. It demonstrates idioms for SEC that model and recognize these scenarios.

3.1 SEC primer

A basic knowledge of SEC's configuration language is required to understand the rules presented below. There are nine basic rule types. I break them into two groups: basic and complex rules. Basic rules types perform actions and do not start an active correlation operation that persists in time. These basic types are described in the SEC man page as:

Suppress: suppress matching input event (used to keep the event from being matched by later rules).

Single: match input event and immediately execute an action that is specified by rule.

Calendar: execute an action at specific times using a cron like syntax.

Complex rules start a multi-part operation that exists for some time after the initial event. The simplest example is a `SingleWithSuppress` rule. It triggers on an event and remains active for some time to suppress further occurrences of the triggering event. A `Pair` rule recognizes a triggering event and initiates a search for a second (paired) event. It reduces two separate but linked events to a single event pair. The complex types are described in the SEC man page as:

SingleWithScript: match input event and depending on the exit value of an external script, execute an action.

SingleWithSuppress: match input event and execute an action immediately, but ignore following matching events for the next T seconds.

Pair: match input event, execute the first action immediately, and ignore following matching events until some other input event arrives (within an optional time window T). On arrival of the second event execute the second action.

PairWithWindow: match input event and wait for T seconds for another input event to arrive. If that event is not observed within a given time window, execute the first action. If the event arrives on time, execute the second action.

SingleWithThreshold: count matching input events during T seconds and if given threshold is exceeded, execute an action and ignore all matching events during rest of the time window.

SingleWith2Thresholds: count matching input events during T1 seconds and if a given threshold is exceeded, execute an action. Now start to count matching events again and if their number per T2 seconds drops below second threshold, execute another action.

SEC rules start with a `type` keyword and continue to the next `type` keyword. In the example rules below, ... are used to take the place of keywords that are not needed for the example, they do not span rules. The order of the keywords is unimportant in a rule definition.

SEC uses Perl regular expressions to parse and recognize events. Data is extracted from events by using subexpressions in the Perl regular expression. The extracted data is assigned to numeric variables $1, $2,...,$N where N is the number of subexpressions in the Perl regular expression. The numeric variable $0 is the entire event. For example, applying the regular expression "`([A-z]*): test number ([0-9]*)`" to the event "`HostOne: test number 34`" will assign $1 the value "`HostOne`", $2 the value "`34`", and $0 will be assigned the entire event line.

Because complex rule types create ongoing correlation operations, a single rule can spawn many active correlation operations. Using the regular expression above, we could have one correlation that counted the number of events for `Host` and another separate correlation that counted events for `HostTwo`. Both counting correlations would be formed from the same rule, but by extracting data from the event the two correlations become separate entities.

This data allows the creation of unique contexts, correlation descriptions and coupled patterns linked to the originating event. We will explore these items in more detail later. Remember that when applying a rule, the regular expression or pattern is always applied first regardless of the ordering of the keywords. As a result, references to $1, $2,...,$N anywhere else in the rule refer to the data extracted by the regular expression.

SEC provides a flow control and data storage mechanism called contexts. As a flow control mechanism, contexts allow rules to influence the application of other rules. Contexts have the following features:

- Contexts are dynamically created and often named using data extracted from an event to make names unique.
- Contexts have a defined lifetime that may be infinite. This lifetime can be increased or decreased as a result of rules or timeouts.
- Multiple contexts can exist at any one time.
- A context can execute actions when its lifetime expires.
- Contexts can be deleted without executing any end-of-lifetime actions.
- Rules (and the correlations they spawn) can use boolean expressions involving contexts to determine if they should apply. Existing contexts return a true value; non-existent contexts return a false value. If the boolean expression is true, the rule will execute, if false the rule will not execute (be suppressed).

In addition to a flow control mechanism, contexts also serve as storage areas for data. This data can be events, parts of events or arbitrary strings. All contexts have an associated data store. In this paper, the word "context" is used for both the flow control entity and its associated data store. When a context is deleted, its associated data store is also deleted. Contexts are most often used to gather related events, for example login and logout events for a user. These contexts can be reported to the system administrator if certain conditions are detected (e.g., the user tried to perform a su during the login session).

The above description might seem to imply that a single context has a single data store; this is not always the case. Multiple contexts can share the same data store using the alias mechanism. This allows events from different streams to be gathered together for reporting or further analysis. The ability to extract data from an event and linking the context by name to that event provides a mechanism for combining multiple event streams into a single context that can be reported. For example, if I extract the process ID 345 from syslog events, I can create a context called: `process_345` and add all of the syslog events with the same PID to that event. If I now link the context `process_346` to the `process_345` context, I can add all of the syslog events with the pid 346 to the same context (data store). So now the `process_345/process_346` context contains all of the syslog events from both processes.

```
type=suppress
desc=ignore non-specific paper problem report \
  since prior events have given us all we need.
ptype=regexp
pattern=.  printer: paper problem$
```

Figure 2: A suppress rule that is used to ignore a noise event sent during a printer error.

In the paper, I use the term session. A session is simply a record of events of interest. In general these events will be stored in one or more contexts. If ssh errors are of interest, a session will record all the ssh events into a context (technically a context data store that may be known by multiple names/aliases) and report that context. If tracing the identities that a user assumes during a login is needed, a different series of data is recorded in a context (data store): the initial ssh connection information is recorded, the login event, the su event as the user tries to go from one user ID to another.

The rest of the elements of SEC rules will be presented as needed by the examples.

3.2 Responding to or filtering single events

The majority of items that we deal with in processing a log file are single items that we have to either discard or act upon. Discardable events are the typical noise where the problem is either not fixable, for example a failing reverse DNS lookups on remote domains from tcp wrappers, or are valueless information that we wish to discard.

Discardable events can be handled using the suppress rule. Figure 2 is an example of such a rule. Since all of my rule sets report anything that is not handled, we want to explicitly ignore all noise lines to prevent them from making it to the default "report everything" rule.

This is a good time to look at the basic anatomy of a SEC rule. All rules start with a `type` option as described in section 3.1. All rules have a `desc` option that documents the rule's purpose. For the complex correlation rules, the description is used differentiate between correlation operations derived from a single rule. We will see an example of this when we look at the horizontal port scan detection rules.

Most rules have a `pattern` option that is applied to the event depending on the `ptype` option. The pattern can be a regular expression, a substring, or a truth value (TRUE or FALSE). The `ptype` option specifies how the `pattern` option is to be interpreted: a regular expression (`regexp`), a substring (`substr`), or a truth value (`TValue`). It also determines if the pattern is successfully applied if it matches the event match (`regexp`/`substr`), or does not match (`nregexp`/`nsubstr`) the event. For `TValue` type patterns, TRUE matches any event (successfull application), while FALSE (not successfully applied) does not match any input event. If the pattern does not successfully apply, the rule is skipped and the next rule in the configuration file is applied.

A number can be added to the end of any of the `nregexp`, `regexp`, `substr`, or `nsubstr` values to make the pattern match across that many lines. So a `ptype` value of `regexp2` would apply the pattern across two lines of input.

By default when an event triggers a rule, the event is *not compared* against other rules in the same file. This can be changed on a per rule basis by using the `continue` option[1].

After single event suppression, the next basic rule type is the single rule. This is used to take action when a particular event is received. Actionable events can interact with other higher level correlation events: adding the event to a storage area (context), changing existing contexts to activate or deactivate other rules, activating a command to deal with the event, or just reporting the event. Figure 3 is an example of a single rule that will generate a warning if the printer is offline from an unknown cause.

In Figure 3 we see two more rule options: `context` and `action`. The `context` option is a boolean expression of contexts that further constrains the rule.

When processing the event "`lj2.cs.umb.edu: printer: Report Printer Offline if needed`", the single rule in Figure 3 checks to see if the

[1]Note: `continue` is not supported for the `suppress` rule type.

```
type=single
continue=dontcont
desc = Report Printer Offline if needed
ptype=regexp
pattern=^([\w._-]+): printer: Report Printer Offline if needed
context = Report_Printer_$1_Offline
action = write - "printer $1 offline, unknown cause" ; \
         delete Report_Printer_$1_Offline
```

Figure 3: A single command that writes a warning message and deletes a context that determines if it should execute.

pattern applies successfully. In this case the pattern matches the event, but if the `Report_Printer_lj2.cs.umb.edu_Offline` context does not exist, then the actions will not be executed. The context `Report_Printer_lj2.cs.umb.edu_Offline` is deleted by other rules in the ruleset (not shown) if a more exact diagnosis of the cause is detected. This suppresses the default (and incorrect) report of the problem.

The `action` option specifies the actions to take when the rule fires. In this case it writes the message `printer lj2.cs.umb.edu offline, unknown cause` to standard output (specified by the file name "`-`") . Then it deletes the context `Report_Printer_lj2.cs.umb.edu_Offline` since it is no longer needed.

There are many potential actions, including:

- creating, deleting, and performing other operations on contexts
- invoking external programs
- piping data or current contexts to external programs
- resetting active correlations
- evaluating Perl mini-programs
- setting and using variables
- creating new events
- running child processes and using the output from the child as a new event stream.

We will discuss and use many of these actions later in this paper.

3.3 Scheduling events with with finer granularity

Part of modeling normal system activity includes accounting for scheduled activities that create events. For example, a scheduled weekly reboot is not worth reporting if the reboot occurs during the scheduled window, however it is worth reporting if it occurs at any other time.

For this we use the `calendar` rule. It allows the reader to schedule and execute actions on a cron like schedule. In place of the `ptype` and `pattern` options it has a `time` option that has five cron-like fields. It is wonderful for executing actions or starting intervals on a minute boundary. Sometimes we need to start intervals with resolution of a second rather than a minute.

Figure 4 shows a mechanism for generating a window that starts 15 seconds after the minute and lasts for 30 seconds. The key is to create two contexts and use both of them in the rules that should be active (or inactive) only during the given window. One context `wait_for_window` expires to begin the timed interval. The `window` context expires marking the end of the interval.

Creating an event on a non-minute boundary is trivial once the reader learns that the event command has a built in delay mechanism. Figure 5 shows the trivial example of generating an event at 30 seconds after the minute and a single rule reacting to the

```
type=calendar
time=30 3 * * *
desc=create 30 second window
action=create window 45; \
       create wait_for_window 15

type=single
...
context=window && !wait_for_window
```

Figure 4: A mechanism for creating an timed interval that starts on a non-minute boundary.

```
type=calendar
time=30 3 * * *
desc=create an event with 30 second delay
action=event 30 Here is the event

type=single
...
pattern=^Here is the event$
...
```

Figure 5: A mechanism for sending an event on a non-minute boundary.

event. Triggering events generated by calendar rules or by expiring contexts can execute actions, define intervals, trigger rules or pass messages between rules. Triggering events are used extensively to detect missing events.

3.4 Detecting missing events

The ability to generate arbitrary events and windows with arbitrary start and stop times is useful when detecting missing events. The rules in Figure 6 report a problem if a '`sendmail -q`' command is not run by root near 31 minutes after the hour. Because of natural variance in the schedule, I expect and accept a sendmail start event from 5 seconds before to 10 seconds after the 31st minute.

The event stream from Figure 1 is used as input to the rules. Figure 7 displays the changes that occur while processing the events in time. Contexts are represented by rectangles, the length of the rectangle is the context's lifetime. Upside down triangles represent the arrival of events. Regular triangles represent actions within SEC. The top graph shows the sequence when the sendmail event fails to arrive, while the bottom graph shows the sequence when the sendmail program is run.

The correlation starts when rule 2 (the `calendar` rule) creating the context `sendmail_31_minute` that will execute an action (write a message to standard output) when it times out after 70 seconds (near 31 minutes and 10 seconds) ending the interval. The calendar rule creates a second context, sendmail_31_minute_inhibit, that will timeout in 55 seconds (near 30 minutes and 55 seconds) starting the 15 second interval for the arrival of the sendmail event. Looking at the top graph in Figure 7, we see the creation of the two contexts on the second and third lines. No event arrives within the 15 second window, so the `sendmail_31_minute` expires and executes the "write" action. The bottom graph shows what happens if the sendmail event is detected. Rule 1 is triggered by the sendmail event occurring in the 15 seconds window and deletes the `sendmail_31_minute` context. The deletion also prevents the "write" action associated with the context from being executed.

Note that the boolean context of rule 1 prevents its execution if the sendmail event were to occur less than 5 seconds before the 31st minute since `! sendmail_31_minute_inhibit` is false because `sendmail_31_minute_inhibit` exists and is therefore true. If the sendmail event occurs after 31 minutes and 10 seconds, the context is again false since `sendmail_31_minute` does not exist, and is false.

The example rules use the write action to report a problem. In a real ruleset, the systems administrator could use the SEC `shellcmd` action to invoke logger(1) to generate a syslog event to be forwarded to a central syslog server. This event would be found by SEC running on the syslog master. The rule matching the event could notify the administrator via email, pager, wall(1) or send a trap to an NMS like HPOV or Nagios. Besides reporting, the event could be further processed with a threshold rule that would try to restart cron as soon as two or more "missed sendmail events" events are reported, and report a

```
# rule 1: detect the sendmail event
type = single
desc = sendmail has run, don't report it as failed
ptype = regexp2
pattern = ^\>  CMD: /usr/lib/sendmail -q.*\n\>  root ([0-9]+) c .*
context = sendmail_31_minute && ! sendmail_31_minute_inhibit
action = delete sendmail_31_minute

# rule 2: define the time window and prep to report a missing event
type = calendar
desc = Start searching for sendmail invocation at 31 past hour
time=30 * * * *
action = create sendmail_31_minute 70 write - \
             Sendmail failed to run detected at %t; \
         create sendmail_31_minute_inhibit 55
```

Figure 6: Rules to detect a missed execution of a sendmail process at the appointed time.

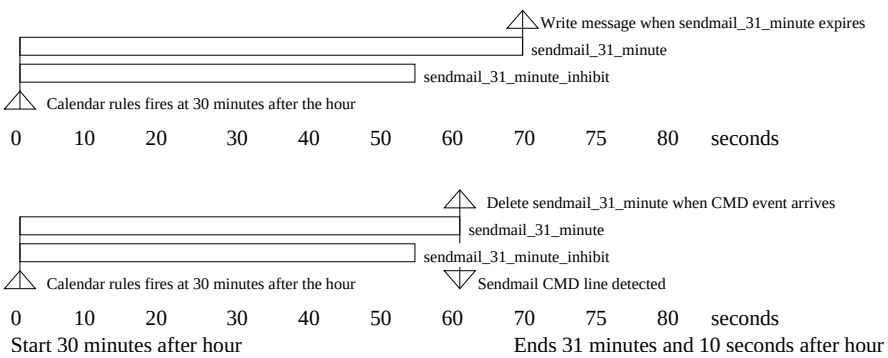

Figure 7: Two timelines showing the events and contexts involved in detecting a missing, or present, sendmail invocation from cron.

problem only if a third consecutive "missed sendmail event" arrived.

Another way to perform missing event detection is to use a pair rule to detect the missing event.

The ruleset of Figure 8 again uses a calendar rule to start the operation. However rather than using contexts, it generates an event at 30:55 to start the 15 second window. This event triggers/arms the `PairWithWindow` rule for 15 seconds. If the sendmail CMD event occurs within the 15 second window of the `PairWithWindow` rule, the `action2` action is executed, which simply logs the successful detection of the sendmail event. If the 15 second window expires without the sendmail event, the `write` action defined by the `action` option is executed.

3.5 Counting correlation

The simplest correlations are counting calculations. A certain number (N) of matching events (E) occur in a given (potentially infinite) time period (T) and are reported as a single composite event (C). This type of correlation reduces reported events by buffering errors until they have reached a critical threshold. Counting correlations have four main variants based on how they reset.

1. The first type of counting correlation uses a sliding window of T seconds. It reuses received events to try to find an interval in which the threshold of N events will be exceeded. The window initially starts at the first occurrence of E. If the time T expires without seeing N items, a new window T starts at the arrival time associated with the next event (E+1) to be seen. The window continues to slide over the arriving events until no more events are available causing the correlation to end. If N items are seen in T seconds, the composite event C is generated and further occurrences of the event E are ignored until the time period T expires.

2. The second type of counting correlation uses a single fixed time period T. The counting operation expires exactly T seconds after the first occurrence of E. This is used when the following events occur in a fixed relationship to the initial event. It is used for tests where the distribution of events within the window is important, for example if the test requires at least two events per minute for the last 10 minutes. This is quite different from saying that there must be 20 events in 10 minutes.

3. The third type of counting correlation resets immediately upon receiving N events within T seconds. It does not suppress the events for the rest of the time window T. When N events are seen, the event C is generated. The rule resets so the next event E starts a new window T. This threshold rule generates 1/N as many events as the original event stream. It keeps generating C events until a time period T has passed. The window T at that point can be sliding or fixed. It is useful for reducing the event flow to upstream correlators while doing a better job of preserving event rate information. This is the closest to a pure counting operation as it generates the event C in a pure ratio to the original incoming events.

4. The last type of counting correlation resets when E's rate of arrival falls below a threshold once the initial threshold has passed. After the first threshold has passed, the correlation continues to suppress the event E until the arrival rate falls below some number N2 in time T2. A version of this type of correlation operation is seen in well-known applications such as HPOV NNM[NNM] where it is referred to as threshold/rearm. This counting operation is useful for software like apache, or logins that have a standard connection/activity rate. Exceeding that rate should trigger a warning that stays active until the event rate drops back to a normal range. Denial of service attacks often exhibit this pattern.

Sadly, the most widely used tool on Unix machines for logging – syslog – causes problems with counting correlations by suppressing duplicate events. We are all familiar with the "Last message repeated N times" entry from syslog that destroys useful information that can be used for threshold modeling. Attempts to work around this problem can generate the proper

```
# rule 1: report a missing event, or log successful detection
type = PairWithWindow
desc = Report failed sendmail
ptype=regexp
pattern= sendmail_31_minute
action = write - Sendmail failed to run detected at %t
desc2 = found sendmail startup entry
ptype2 = regexp2
pattern2 = ^\>  CMD: /usr/lib/sendmail -q.*\n\>  root ([0-9]+) c .*
action2 = logonly
window=15

# rule 2: trigger the detection operation at 30:55.
type = calendar
desc = Start searching for sendmail invocation at 31 past hour
time=30 * * * *
action = event 55 sendmail_31_minute
```

Figure 8: A rule to detect a failed execution of a sendmail process at the appointed time using the `PairWithWindow` rule.

number of events, but cannot recover the inter-event timing information between the duplicate events that is lost forever.

3.5.1 Implementing a cooling-off period

One problem with counting is that there may be an initial flurry of events that stops automatically as the problem is diagnosed and solved by the person or device that caused it. I see this with "disk quota exceeded" events. A flurry of alerts is generated immediately as the user's program tries to write to the filesystem. Depending on the experience level of the user, this may occur for a minute or so. I wish to receive an alert when there have been more than three quota exceeded messages for the same user/machine/filesystem triplet. These three events must occur in a fixed 2-minute window that starts 1 minute after the first quota exceeded message occurs. This counting operation is of the second type explained above.

If after this amount of time, the user is still experiencing a problem, then they may not realize what is happening or be able to deal with the problem. In any case, they can use some assistance.

A simple counting rule will include the initial one minute period where I expect a flurry of activity. This makes it difficult to set thresholds because of the wide variance in the number of events that can occur. By deferring the count to a sampling range that has a lower variance, I am able to reduce the number of false positives.

Figure 9 shows the rules that I use to implement this type of operation, and Figure 10 shows the time graph of two event streams using these three rules. In Figure 10 we add a new item to the graph. The rectangular contexts, upside down event triangles and regular action triangles are present, but we also have rectangles with triangular ends. These 6 sided figures represent active correlation operations. The length of the hextangle represents the lifetime of the correlation operation. We see two separate event streams in Figure 10. Both streams come from the same user on the same host, but occur on different filesystems. The E1 stream is caused by a quota on the /home filesystem, while the E2 stream is on the /var filesystem.

Starting at the left hand side of Figure 10, we see the first event of the E1 stream arrive. It triggers rule 1 of Figure 9 and causes the

```
# Example Input:
# Jul 11 11:50:38 host1.example.org ufs: [ID 725924 \
#      kern.notice] quota_ufs: over hard disk limit \
#      (pid 0, uid 4816, inum 505699, fs /mount/sd0f)
# rule 1: perform the count after the one minute interval
type=SingleWithThreshold
desc=Detect user pid $2 on $1 (fs $3) quota issue
continue=takenext
ptype=regexp
pattern=([\w._-]+) ufs: \[.*uid (\d+).*, fs ([\w/._-]+)
context= ! wait_a_minute_$1_$2_$3
action= write - Quota issue $0;
thresh=4
window=180

# rule 2:  define the beginning of the one minute interval
#          and define a fixed three minute counting window.
type=single
desc=Wait one minute for $1 $2 $3
ptype=regexp
pattern=([\w._-]+) ufs: \[.*uid (\d+).*, fs ([\w/._-]+)
context= ! in_threshold_for_$1_$2_$3
action=create in_threshold_for_$1_$2_$3 180 \
              reset Detect user pid $2 on $1 (fs $3) quota issue ;\
       create wait_a_minute_$1_$2_$3 60

# rule 3: discard events during the three minute correlation interval
type=suppress
desc=Trap events
ptype=regexp
pattern=([\w._-]+) ufs: \[.*uid (\d+).*, fs ([\w/._-]+)
```

Figure 9: Rules to handle the flurry of activity that occurs when exceeding a quota. Reports when more than 3 events received in two minutes after a one minute delay.

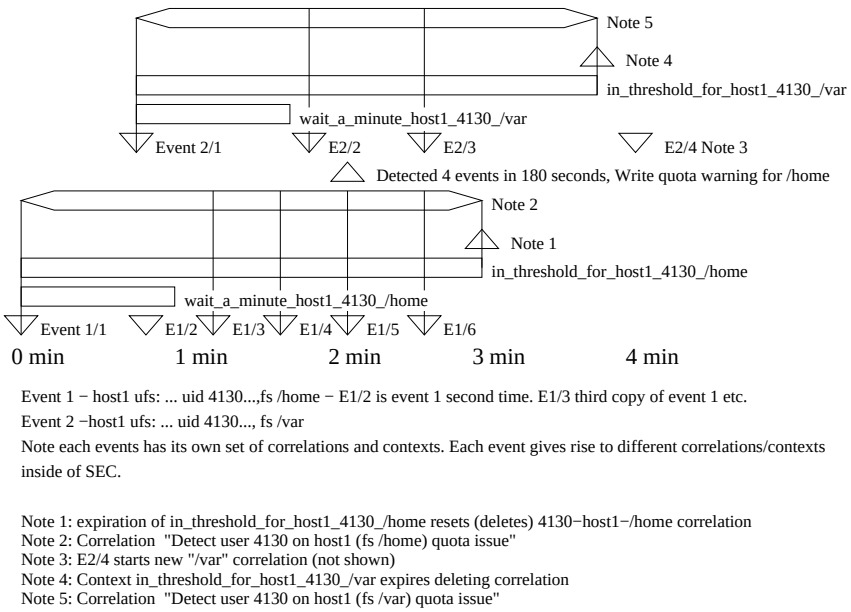

Event 1 − host1 ufs: ... uid 4130...,fs /home − E1/2 is event 1 second time. E1/3 third copy of event 1 etc.
Event 2 −host1 ufs: ... uid 4130..., fs /var
Note each events has its own set of correlations and contexts. Each event gives rise to different correlations/contexts inside of SEC.

Note 1: expiration of in_threshold_for_host1_4130_/home resets (deletes) 4130−host1−/home correlation
Note 2: Correlation "Detect user 4130 on host1 (fs /home) quota issue"
Note 3: E2/4 starts new "/var" correlation (not shown)
Note 4: Context in_threshold_for_host1_4130_/var expires deleting correlation
Note 5: Correlation "Detect user 4130 on host1 (fs /var) quota issue"

Figure 10: Time graph of two event streams E1 and E2 using the rules of Figure 9.

creation of the `Detect user 4130 on host1 (fs /home) quota issue` correlation represented on line 5 of Figure 10. It will trigger when it counts three more events (for a total of 4 events) in the next three minutes. Because rule 1 has the `continue` option set to `takenext`, the E1/1 event also triggers rule 2. Rule 2 creates the context `in_threshold_for_host1_4130_/home` for 180 seconds. It also creates the `wait_a_minute_host1_4130_/home` for 60 seconds.

The E1/2 event arrives while the `wait_a_minute_host1_4130_/home` still exists, so rule 1's context prevents the event from being counted (note the lack of a vertical line), and the existence of the `in_threshold_for_host1_4130_/home` context prevents rule 2 from firing. So rule 3 fires suppressing the E1/2 event.

When the E1/3 event occurs, the `wait_a_minute_host1_4130_/home` has expired, so the event is counted by the `Detect user 4130 on host1 (fs /home) quota issue` correlation operation on line 5. Due to the `continue` setting, the event is passed to rule 2 that ignores it, since the `in_threshold_for_host1_4130_/home` context is still active. So rule 3 consumes the event. The E1/4 and E1/5 events follow the same pattern as the E1/3 event, and causes the correlation to reach its threshold of 4 events. At this point the quota warning is written as shown by the triangle on line 6 of Figure 10. Event E1/6 is ignored by rule 1 since the threshold has been exceeded. The event is then passed to rule 2 and ignored by the context. The event is finally consumed by rule 3.

After three minutes, the `in_threshold_for_host1_4130_/home` context expires causing its action to fire, which resets the `Detect user 4130 on host1 (fs /home) quota issue` correlation. The next time an event from the E1 stream occurs, it will cause a new chain of correlations and counting.

Note that while we were busy with the E1 stream, the E2 stream was also being processed. However only three events arrived before the `in_threshold_for_host1_4130_/var` context expired deleting the `Detect user 4130 on host1 (fs /var) quota issue` context before it reached the threshold to execute its action. It is important to

note that two different and separate correlation operations were in effect with both of them using the same rules. The dynamic extraction of data from the events allowed the creation of these two unique correlation operations even though they were spawned from the same series of three rules.

3.5.2 Reporting too few events in a time period

Sometimes we want to take action when too few events have occurred. Figure 11 demonstrates a ruleset that sends an alert if less than 50 HTTP GET requests arrive in a hour. This problem can have multiple causes, including DNS failure, abnormal server exit, and network connectivity problems.

Figure 11 combines a `SingleWithThreshold` rule with a single rule that are both triggered by the same event. The single rule (rule 1) creates a context that:

- generates an event to to report the problem.
- resets the `SingleWithThreshold` rule (rule 2) so that the time periods for the `too_few_http_accesses` context and counting correlation are synchronized.
- generates a GET event to start a new correlation window and initialize a new `too_few_http_accesses` context.

Rule 1 creates a context that generates an event to trigger a report rather than using the `write` action to report a problem directly. Generating a report event allows the suppression of the report by another rule (rule 3) based upon other factors (e.g., at 2AM, 50 gets/hour is not realistic.)

This ruleset is similar to other rules that detect a missing event. In this case the "event" we are missing is a sufficient number of GET requests. In this example a rule 4, a calendar rule, is used to prime the ruleset so that its window starts on the hour. Every hour we reset the correlation operation created by rule 2 and cause the expiration of the `too_few_http_accesses` context using the obsolete action. If the `too_few_http_accesses` context exists, then it will execute its action list. If the context does not exist, then nothing will happen. Lastly, rule 4 generates a new event that triggers (primes) rules 1 and 2.

Another variant of this rule does not use a calendar rule to start counting at a regular interval. It starts counting after another event has been seen. For example, you can make sure that all the embedded images on a web page are transferred by using the get of the web page as the triggering event, and then counting the number of files transferred to the IP address in the following N minutes. This creates cross event counting since the single rule does not have to have the same triggering event as the `SingleWithThreshold` rule.

3.6 Repeat Elimination/Compression

I have dealt with real-time log file reporters that generated 300 emails when a partition filled up overnight. Thus, there must be a method to condense, deduplicate or eliminate repeated events to provide a better picture of a problem, and reduce the number of messages to prevent the monitoring system from spamming the administrators.

The `SingleWithSuppress` rule fills this deduplication need. To handle file system full errors, the rule in Figure 12 is used.

This rule reports that the filesystem is full when it receives its first event. It then suppresses the event message for the next hour. Note that the `desc` keyword includes the filesystem and hostname ($2 and $1 respectively). This makes the correlation operation that is generated from the rule unique so that a disk full condition on the same host for the filesystem /mount/fs2 will generate an error event if it occurs 5 minutes after the /mount/sd0f event. If the filesystem was not included in the `desc` option, then only one alert for a full filesystem would be generated regardless of how many filesystems actually filled up during the hour.

Another way of performing duplicate elimination is to use a `SingleWith2Thresholds` rule. Figure 13 shows handling the same error using this rule. The `SingleWith2Thresholds` rule does not reset on a time basis like `SingleWithSuppress`. It resets when the number of events coming in drops below the sec-

```
# rule 1
type = single
continue = takenext
desc = report not enough http requests
ptype = substr
pattern = GET
context = ! too_few_http_accesses
action = create too_few_http_accesses 3600 \
                (event 0 REPORT_TOO_FEW_HTTP_ACCESSES ; \
                 reset  Look for enough http accesses ; \
                 event 0 GET )

# rule 2 - count the number of events and delete report
#          context if enough events have occurred.
type = SingleWithThreshold
desc = Look for enough http accesses
ptype=substr
pattern = GET
action = delete too_few_http_accesses
thresh = 50
window= 3600

# rule 3 - generate the report if we are in the correct time
#          window
type= single
desc = report too few get events
...
pattern = ^REPORT_TOO_FEW_HTTP_ACCESSES$
context = ! inhibit_low_http_report_between_midnight_and_7am
action = write - low http get count

# rule 4 - A priming rule to start measurements every hour on the hour.
type=calendar
desc = trigger ruleset looking for too few http requests
...
time=0 * * * *
action= obsolete too_few_http_accesses ; \
        reset Look for enough http accesses ; \
        event GET
```

Figure 11: A rule to detect too few HTTP get events in a time period. The obsolete action causes the context to time out, thus running its actions.

```
# Example:
#   Apr 13 15:08:52 host4.example.org ufs: [ID 845546 \
#       kern.notice] NOTICE: alloc: /mount/sd0f: file system full
type=SingleWithSuppress
desc=Full filesystem $2 on $1
ptype=regexp
pattern=([\w._-]+) ufs: \[.* NOTICE: alloc: ([\w/._-]+): file system full
action= write - filesystem $2 on host $1 full
window=3600
```

Figure 12: A rule to report a file system full error and suppress further errors for 60 minutes.

```
# Example:
#   Apr 13 15:08:52 host4.example.org ufs: [ID 845546 \
#       kern.notice] NOTICE: alloc: /mount/sd0f: file system full
type=SingleWith2Thresholds
desc=Detect full filesystem $2 on $1
ptype=regexp
pattern=([\w._-]+) ufs: \[.* NOTICE: alloc: ([\w/._-]+): file system full
action= write - filesystem $2 on host $1 full
thresh=10
window=60
desc2=Detect end of full filesystem $2 on $1 condition
action2= write - full filesystem $2 on host $1 is resolved
thresh2=1
window2=3600
```

Figure 13: A rule to report a file system full error when more than 10 events are received in 1 minute. Further problems reports are suppressed until the rule resets. The rule resets when no events (less than 1 event) are received in 60 minutes.

ond threshold. In Figure 13 the alert is sent when more than ten "disk full" events occur in a minute. The rule considers the problem to have disappeared if there are no events in one hour. This one hour period could be at 8AM on the Tuesday morning after a long weekend. If the filesystem filled up Friday night at 9PM, and the "disk full" event was reported via email, the sysadmin will most likely only receive one email, assuming that at least one file system full message is generated every 60 minutes. Using the `SingleWithSuppress` rule, the administrator will have 59 emails, one for every one hour period. The worst case for the two thresholds rule is no worse than the `SingleWithSuppress` implementation, and may be better.

Looking at this, the reader must wonder when to use `SingleWithSuppress`. It is useful when the status of a device needs constant refreshing. For example, sending traps to a network management station indicates that the problem is still ongoing. `SingleWithSuppress` can be combined with other correlation rules by generating an event when the rule fires. This generated event can be thresholded to perform more complex counting actions. For example, if a `SingleWithSuppress` rule detects and suppresses NFS errors for an hour, the staff can be alerted if it has triggered once an hour for the last 5 hours.

3.7 Report on analysis of event contents

Unlike most other programs, SEC allows the reader to extract and analyze data contained within an event. One simple example is the rule that analyzes NTP time adjustments. Many of the systems at UMB are housed in offices rather than in the computer room. Temperature changes and other factors contribute to clock drift. With the frequency of time adjustments, I consider any clock with less than 1/4 a second difference from the NTP controlled time sources to be within a normal range. Figure 14 shows the rules that are applied to analyze the xntpd time adjustment events. We extract the value of the time change from the step messages. This value is assigned to the variable $1. The context expression executes a Perl mini-program to see if the absolute value of the change is larger than the threshold of 0.25 seconds. If it is, the context is satisfied and the rule's actions fire.

The context expression uses a mechanism to run arbitrary Perl code. It then uses the result of the expression to determine if the rule should fire. It can be used to match networks after applying a netmask, perform calculations with fields of the event or other tasks to properly analyze the events.

Although this idea is illustrated using a `Single` rule, this particular tests would be better implemented using a `SingleWithThreshold` rule. That way an alert is generated only if a certain number of time changes in the past 6 hours failed, the administrator can check the machine, make sure the air conditioner is working properly etc. This also allows multiple parameters, for example, changes greater than 10 seconds are handled by a single rule that immediately notifies the admins, while smaller changes are handled by a threshold rule that allows for some variance in the readings.

3.8 Detect identical events occurring across multiple hosts

A single incident can affect multiple hosts. Detecting a series of identical events on multiple hosts provides a measure of the scope of the problem. The problem can be an NFS server failure affecting only one host that does not need to be paged out in the middle of the night, or it may affect 100 hosts, which requires recovery procedures to occur immediately. Other problems such as time synchronization, or detection of port scans also fall into this realm.

One typical example of this rule is to detect horizontal port scans. The rules in Figure 15 identify a horizontal port scan as three or more connection denied events from different server hosts within 5 minutes from a particular external host or network. So 20 connections to different ports on the same host would not result in the detection of a horizontal scan. In the example, I assume that the hosts are equipped with TCP wrappers that report denied connections. The set of rules in Figure 15 implements the detection of a horizontal port scan by counting unique client host/server host combinations.

```
type=single
desc = report large xntpd corrections for host $1
continue = dontcont
ptype=regexp
context= =(abs($2) > 0.25)
pattern=([A-z0-9._-]+) xntpd\[[0-9]+\]:.*time reset \(step\) ([-]?[0-9.]+) s
action= write - "large xntpd correction($2) on $1"
```

Figure 14: Rule to analyze time corrections in NTP time adjustment events. The absolute value of the time adjustment must be greater than 0.25 seconds to generate a warning.

A timeline of these three rules is shown in Figure 16.

The key to understanding these rules is to realize that the description field is used to match events with correlation operations. When rule 1, the threshold correlation rule, sees the first rejected connection from 192.168.1.1 to 10.1.2.3, it generates a `Count denied events from 192.168.1.1` correlation. The next time a deny for 192.168.1.1 arrives, it will be tested by rule 1, the description field generated from this new event will match an ongoing correlation threshold operation and it will be considered part of the `Count denied events from 192.168.1.1` threshold correlation. If a rejection event for the source 193.1.1.1 arrives, the generated description field will not match an active threshold correlation, so a new correlation operation will be started with the description `Count denied events from source2`. Figure 16 shows a correlation operation from start to finish. First the event E1 reports a denial from host 192.168.1.1 to connect/scan 10.1.2.3. The correlation operation `Count denied events from 192.168.1.1` is started by rule 1, rule 2 is skipped because the pattern does not match, and rule 3 creates the 5-minute-long context `seen_connection_from_192.168.1.1_to_10.1.2.3` that is used to filter arriving event to make sure that only unique events are counted. The rest of rule 3's actions will be discussed later.

The count for rule 1, the threshold correlation operation, is incremented only if the `seen_connection_from_192.168.1.1_to_10.1.2.3` context does not exist. When the E1/2 (event 1 number 2) arrives, this context still exists and all the rules ignore the event. When E2/1 arrives, it triggers rule 1 and rule 3 creating the appropriate context and incrementing the threshold operation's count.

When five minutes have passed since E1/1's arrival and the threshold rule has not been triggered by the arrival of three events, the start of the threshold rule is moved to the second event that it counted, and the count is decremented by 1. This occurs because the threshold rule uses a sliding window by default. When events 3/1 and 4/1 arrive, they are counted by the shifted threshold correlation operation started by rule 1. With the arrival of E2/1, E3/1, and E4/1, three events have occurred within five minutes and a horizontal port scan is detected. As a result, the action reporting the context `conn_deny_from_192.168.1.1` is executed and the events counted during the correlation operation (maintained by the add action of rule 3) are reported to the file report_log.

Rule 2 and the final actions of rule 3 allow detection of horizontal port scans even if they come from different hosts such as: 192.168.3.1, 192.168.1.1, and 192.168.7.1. If each of these hosts scans a different host on the 10 network, it will be detected as a horizontal scan from the 192.168.0.0 network. This is done by creating three events replacing the real source address with a corresponding network address. One event is created for each class A, B and C network that the original host could belong to: 192.168.1.0, 192.168.0.0, and 192.0.0.0. The response to these synthesized events are not shown in Figure 16, but they start a parallel series of correlation

```
# Example input:
# May 10 13:52:13 cyber TCPD-Event cyber:127.6.7.1:3424:sshd deny \
#     badguy.example.com:192.268.15.45 user unknown
#  Variable = description (value from example above)
#  $3  = server ip address (127.6.7.1)
#  $5  = daemon or service connected to on server (sshd)
#  $8  = ip address of client (attacking) machine (192.268.15.45)
#  $9  = 1st quad of client host ip address (192)
#  $10  = 2nd quad of client host ip address (6)
#  $11 = 3rd quad of client host ip address (7)
#  $12 = 4th quad of client host ip address (1)
# Rule 1: Perform the counting of unique destinations by client host/net
type = SingleWithThreshold
desc = Count denied events from $8
continue = takenext
ptype = regexp
pattern = ^(.*) TCPD-Event ([A-z0-9_.]*):([0-9.]*):([0-9]*):([^ ]*) (deny) \
            ([^:]*):(([0-9]*)\.([0-9]*)\.([0-9]*)\.([0-9]*)) user (.*)
action = report conn_deny_from_$8 /bin/cat >> report_log
context = ! seen_connection_from_$8_to_$3
thresh = 3
window = 300

## Rule 2: Insert a rule to capture synthesized network tcpd events.
type=single
...
pattern = ^(.*) TCPD-Event ([A-z0-9_.]*):([0-9.]*):([0-9]*):([^ ]*) (deny) \
            ([^:]*):(([0-9]*)\.([0-9]*)\.([0-9]*)\.([0-9]*)) user (.*) net$
action=none

# Rule 3: Generate network counting rules and maintain contexts
type = single
desc = maintain counting contexts for deny service $5 from $8 event
continue = takenext
ptype = regexp
pattern = ^(.*) TCPD-Event ([A-z0-9_.]*):([0-9.]*):([0-9]*):([^ ]*) (deny) \
            ([^:]*):(([0-9]*)\.([0-9]*)\.([0-9]*)\.([0-9]*)) user (.*)
context = ! seen_connection_from_$8_to_$3
action = create seen_connection_from_$8_to_$3 300; \
         add conn_deny_from_$8 $0 ; \
 event 0 $1 TCPD-Event $2:$3:$4:$5 $6 $7:$9.$10.$11.0 user $13 net; \
 event 0 $1 TCPD-Event $2:$3:$4:$5 $6 $7:$9.$10.0.0 user $13 net; \
 event 0 $1 TCPD-Event $2:$3:$4:$5 $6 $7:$9.0.0.0 user $13 net; \
         add conn_deny_from_$9.$10.$11.0 $0 ; \
         add conn_deny_from_$9.$10.0.0 $0 ; \
         add conn_deny_from_$9.0.0.0 $0
```

Figure 15: Rules to detect horizontal port scans defined by connections to 3 different server hosts from the same client host within 5 minutes. Note: patterns are split for readability. This is not valid for sec input.

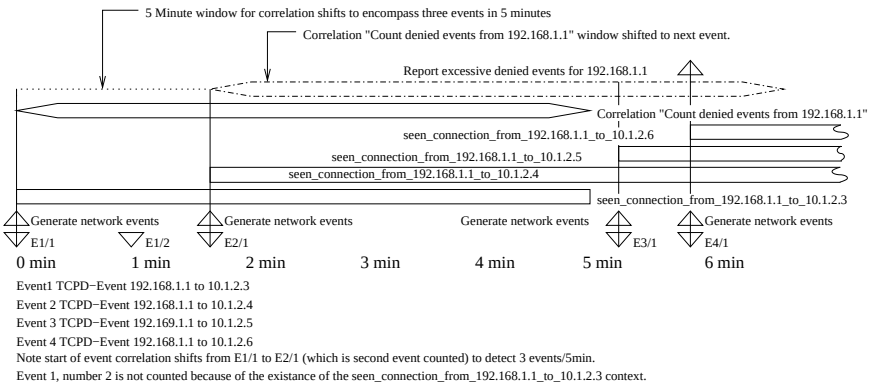

Figure 16: Timeline showing the application of rules to detect horizontal port scans.

operations and contexts using the network address of the client in place of 192.168.1.1.

Vertical scans can use the same framework with the following changes:

- the filtering context needs to include port numbers so that only unique client host/server host/port triples are counted by the threshold rule.
- the description of rule 1 to include the server host IP so that it only counts connections to a specific server host.

This will count the number of unique server ports that are accessed on the server from the client host.

In general, using rules 1 and 3, you can count unique occurrences of a value or group of values. The context used to link the rules must include the unique values in its name. The description used in rule 1 will not include the unique values and will create a bucket in which the events will be counted. In the horizontal port scan case, case, my bucket was any connection from the same client host. The unique value was the server IP address connected to by the the client host. In detecting a vertical port scan, the value is the number of unique ports connected to while the bucket is the client/server host pair.

These two changes allow the counting ruleset to count the number of unique occurrences of the parameter that is present in the filtering rule, but missing from the rule 1 description (the bucket). E.g., if the context specifies serverhost, clienthost, serverport and rule 1 specifies clienthost and serverhost in its description, then the rules above implement counting of unique ports for a given clienthost and serverhost. The rules as presented above specified clienthost and serverhost, rule 1 specified the clienthost, so the ruleset counted unique serverhost's for a given clienthost.

Other counting methods can also be implemented using mixtures of the vertical and horizontal counting methods.

While I implemented a "pure" SEC solution, the ability to use Perl functions and data structured from SEC rules provides other solutions[Vaarandi7_2003] to this problem.

3.9 Creating threads of events from multiple sources

Many thread recognition operations involve using one of the pair type rules. Pair rules allow identification of a future (child) event by searching for identifying information taken from the present (parent) event. This provides the ability to stitch a thread through various events by providing a series of pair rules.

There are three times when you need to trigger an action with pair rules:

1. Take action upon receipt of the parent event
2. Take action upon receipt of the child event

3. Take action after some time when the child event has not been received (expiration of the pair rule).

The `Pair` rule provides actions for triggers 1 and 2. The `PairWithWindow` rule provides actions for triggers 2 and 3. None of the currently existing pair rules provides a mechanism for taking actions on all three triggers. Figure 17 shows a way to make up for this limitation by using a context that expires when the pair rule is due to be deleted. Since triggers 2 and trigger 3 are mutually exclusive, part of trigger 2's action is to delete the context that implements the action for to trigger three.

I have used this method for triggering an automatic repair action upon receipt of the first event. The arrival of the second event indicated that the repair worked. If the second event failed to arrive, an alert would be sent when the context timed out. Also, I have triggered additional data gathering scripts from the first event. The second event in this case reported the event and additional data when the the end of the additional data was seen. If the additional data did not arrive on time, I wanted the event to be reported.

This mechanism can replace combinations of `PairWithWindow` and `Single` rules. It simplifies the rules by eliminating duplicate information, such as patterns, that must be kept up to date in both rules.

3.9.1 Correlating across processes

One of more difficult correlation tasks involves creating a session made up of events from multiple processes.

Figure 18 shows a ruleset that sets up a link between parent and child ssh processes. Its application is show in Figure 19.

When a connection to ssh occurs, the parent process, running as root, reports the authentication events and generates information about a user's login. After the authentication process, a child sshd is spawned that is responsible for other operations including port forwarding and logout (disconnect) events. The ruleset in Figure 18 captures all of the events generated by the parent or child ssh process. This includes errors generated by the parent and child ssh processes.

A session starts with the initial network connection to the parent sshd and ends with a connection closed event from the child sshd. I accumulate all events from both processes into a single context. I also have rules (not shown in the example) to report the entire context when unexpected events occur.

The tricky part is accumulating the events from both processes into a single context. The connection between the event streams is provided by a tie event that encompasses unique identifying elements from both event streams and thus ties together the two streams into a single stream.

Each ssh process has its own unique event stream stored in the context `session_log_<hostname>_<pid>`. There is a `Single` rule, omitted for brevity, that accumulates ssh events into this context. When the tie event is seen, it provides the link between the parent sshd session_log context and the child session_log context. The data from the two contexts is merged and the two context names (with the parent and child pid's) are assigned to the same underlying context. Hence the child's `session_log_<hostname>_<child pid>` context and the parent's `session_log_<hostname>_<parent pid>` contexts refer to the same data. After the contexts are linked, actions using either the child context name or the parent context name operate on the same underlying context. Reporting or adding to the context using one of the linked names acts the same regardless of which name is used.

In Figure 19 the first event E1 triggers rule 1 from Figure 18, the `PairWithWindow` rule, to recognize the start of the session. The second half of rule 1 looks for a tie event for the following 60 seconds. There may be many tie events, but there should be only one tie event that contains the the pid of the parent sshd. Since we have that stored in $2, we use it in pattern2. The start of session event is passed onto additional rules (not shown) by setting the continue option on rule 1 to `takenext`. These additional rules record the events in the session_log context identified by system and pid, as in the `session_log_example.org_10240` context of Figure 19.

If the tie event is not found within 60 seconds, the `session_log_example.org_10240` context is reported. However, if the tie event is found as in Fig-

```
type=pair
...
action = write - rule triggered ;\
         create take_action_on_pair_expiration 60 (write - rule expired)
...
pattern2=
action2 = write - pattern 2 seen ;\
          delete take_action_on_pair_expiration
...
window=60
```

Figure 17: A method to take an action on all three trigger points in a pair rule.

```
# rule 1 - recognize the start if an ssh session,
#          and link parent and child event contexts.
type=PairWithWindow
continue=takenext
desc=Recognize ssh session start for $1[$2]
ptype=regexp
pattern=([A-Za-z0-9._-]+) sshd\[([0-9]+)\]: \[[^]]+\] Connection from ([0-9.]+) port [0-9]+
action=report session_log_$1_$2 /bin/cat
desc2=Link parent and child contexts
ptype2=regexp
pattern2=([A-Za-z0-9._-]+) [A-z0-9]+\[[0-9]+\]: \[[^]]+\] SSHD child process +([0-9]+) spawned by $2
action2=copy session_log_$1_$2 %b; \
        delete session_log_$1_$2; \
        alias session_log_$1_%2 session_log_$1_$2; \
        add session_log_$1_$2 $0; \
        event 0 %b; \
alias session_log_owner_$1_%2 session_log_owner_$1_$2; \
window=60

# rule 2 - recognize login event and save username for later use
type=single
desc=Start login timer
ptype=regexp
pattern=([A-Za-z0-9._-]+) sshd\[([0-9]+)\]: \[[^]]+\] Accepted (publickey|password) for ([A-z0-9_-]+) from [0-9.]+ port [0-9]+ (.*)
action=add session_log_$1_$2 $0; add session_log_owner_$1_$2 $4

# rule 3 - handle logout
type=single
desc=Recognize ssh session end
ptype=regexp
pattern=([A-Za-z0-9._-]+) sshd\[([0-9]+)\]: \[[^]]+\] Closing connection to ([0-9.]+)
action= delete session_log_$1_$2; delete session_log_owner_$1_$2
```

Figure 18: Accumulating output from ssh into a single context.

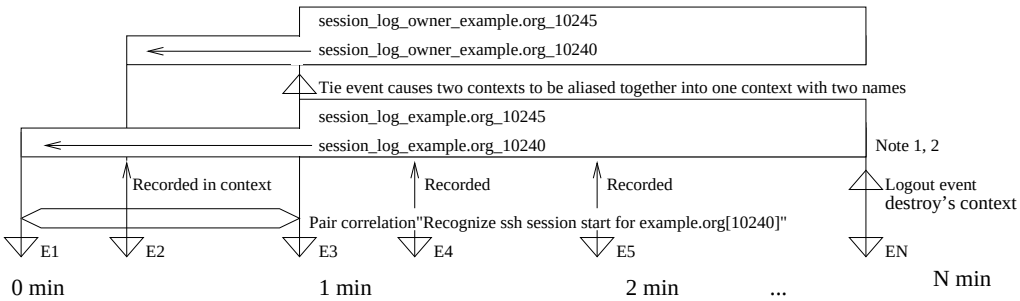

E1 − Initial parent event: example.org sshd[10240]: [ID 800047 auth.info] Connection from 192.168.0.1 port 3500
E2 − Parent event: example.org sshd[10240]: [ID 800047 auth.info] Accepted publickey for rouilj from 192.168.0.1 port 3500 ssh2
E3 − Tie event: example.org rouilj[10248]: [ID 702911 auth.notice] SSHD child process 10245 spawned by 10240
E4 − Child event: example.org sshd[10245]: [ID 800047 auth.info] error: connect_to 127.0.0.1 port2401 failed
E5 − Child event: example.org sshd[10245]: [ID 800047 auth.info] bind: Cannot assign requested address
EN − Logout child event: example.org sshd[10245]: [ID 800047auth.info] Closing connection to 192.168.0.1

Note 1: creation of session_log_example.org_10240 context in response to E1 is done by a catchall rule that is not shown in the ruleset.
Note 2: alias of two contexts is shown by the large box labeled with both context names.

Figure 19: The application of the ssh ruleset showing the key events in establishing the link between parent and child processes.

ure 19, then a number of other operations occur. The tie event is generated by a script that is run by the child sshd. Therefore it is possible for the child sshd to generate events before the tie event is created. Because of the default rule that adds events to the `session_log_example.org_10245`, additional work must be done when the tie event arrives to preserve the data in the child's session_log. The second part of rule 1 in Figure 18 copies child's session_log context into the variable %b. The child's session_log is then deleted and aliased to the parent session_log. The %2 variable is the value of $2 from the first pattern, the parent process's PID. After `pattern2` is applied, the parent PID is referenced as %2 because $2 is now the second subexpression of `pattern2`. Next the data copied from the child log is injected into the event stream to allow re-analysis and reporting using the combined parent and child context.

The last action for the tie event is to alias the login username stored in the context `session_log_owner_<hostname>_<parent pid>` to a similar context under the child pid. Then any rule that analyzes a child event can obtain the login name by referencing the alias context. Rule 2 in Figure 18 handles the login event and creates the context `session_log_owner_<hostname>_<parent pid>` where it stores the login name for use by the other rules in the ruleset. Rule 2 also stores the login event in the session_log context.

The last rule is very simple. It detects the "close connection" (logout) event and deletes the contexts created during the session. The delivery of event N (EN) in Figure 19 causes deletion of contexts. Deleting an aliased contexts deletes the context data store as well as all the names pointing to the context data store. Rule 3 uses the child PID to delete `session_log_example.org_10245` and `session_log_owner_example.org_10245`, which cleans up all 4 context names (2 from the parent PID and 2 from the child) and both context data stores.

This mechanism can be used for correlating any series of events and passing information between the rules that comprise an analysis mechanism. The trick is to find suitable tie events to allow the thread to be followed. The tie event must contains unique elements found in the events streams that are to be tied together. In the ssh correlation I create a tie event using the pid's of the parent and child events. Every child event includes the PID of the child sshd so that I can easily construct the context name that

points to the combined context data store. For the ssh correlation, I create the tie event by running shell commands using the sshrc mechanism and use the logger(1) command to inject the tie event into the data stream. This creates the possibility that the tie event arrives after events from the child process. It would make the correlation easier if I modified the sshd code to provide this tie event since this would generate the events in the correct order for correlation.

Having the events arriving in the wrong order for cross correlation is a problem that is not easily remedied. I suppress reporting of the child events while waiting for the tie event (not shown). Then once the tie event is received, the child events are resubmitted for correlation. This is troublesome and error prone and is an area that warrants further investigation.

3.9.2 Recognizing coinciding events

The pair rules by their nature force events to be ordered in time. The parent event must occur before the child event. However, we often need to find out if a group of events arrives in some window. For example, an alert must be generated if any two interfaces on a system have reported an interface down event within 5 minutes of each other. Figure 20 shows a method of testing for this situation using a series of single rules. Each single rule recognizes one condition. It sets a context that exists for 5 minutes recognizing that condition. It then generates an event that causes the last single rule to see if any two conditions are still asserted. Since the time periods for each created context can be different, it allows any type of overlap between the coinciding events. To detect an exact ordering for some of the events, use a pair rule to create an "event_a_occurred_within_5_minutes_after_event_b" context. Then use this context to detect the coincidence of this strict ordering with other events.

```
type=single
...
pattern=interface eri0 down
action = create eri_0_down 300 ; \
         event 0 check_for_eri_down

type=single
...
pattern=interface eri1 down
action = create eri_1_down 300 ; \
         event 0 check_for_eri_down

type=single
...
pattern=interface eri2 down
action = create eri_2_down 300 ; \
         event 0 check_for_eri_down

type=single
...
pattern= check_for_eri_down
context = ( eri_1_down && eri_2_down ) || \
          ( eri_1_down && eri_3_down ) || \
          ( eri_2_down && eri_3_down )
action=...
```

Figure 20: Identify combinations of two events occurring within a 5 minute window

3.9.3 Handling a sequence

I refer to a known explicitly ordered list of events from a single application or device as a sequence. Much of normal operation is expressed as a sequence. An

example is the sequence of events generated for a rebooting system. There are many places where things can go wrong, from failing hardware exposed by the stress of rebooting, to failures caused by software and configuration changes that are only tested on reboot. In these types of sequences, there are a multitude of factors to consider:

- What are the start and end events for the sequence? Examples include login and logout events, or start of reboot to end of reboot phase of a machine.
- Must the end event occur within a specific time range after the start event? Examples include email receipt for a local user to time of final delivery must take less then 1 minute, or the start of backup cron job to end of cron job must take longer than 10 minutes but less than 2 hours.
- Are there inter-event timing constraints? During a system reboot, it should not take more than 10 seconds from the identification of the network interface card till it is configured.
- What information from the sequence is needed to aid diagnosis when a problem occurs? Do we need the entire sequence or do we need to report a fragment of the entire sequence?
- Are there known failure modes for the sequence that require reporting?
- Are there optional items in the sequence that need to be checked?

Setting up rules for a full sequence like a system reboot consisting of 80 lines can be a pain. But, failure to fully specify the sequence can result in false negatives. By blending verification of correct operation and looking for known failure modes, the risk of false negatives can be reduced. However, as was stated by someone who shall remain anonymous:

> we have no crystal ball to let us know there is a new kernel error message saying 'the CPU is melting!'.

A tool that would take a sequence, analyze it for time interdependencies and generate a ruleset to detect the sequence would be welcome.

Figure 21 is a partial example of the boot sequence for a host[2]. It makes sure that all of the identified components are present and places an overall timing constraint (300 seconds) on the time from the first line to the last in the reboot sequence.

This demonstrates the use of a few things that we have discussed previously. Rule 1 recognizes the start/end of reboot using a three-trigger pair rule. A context is created (`look_for_line_2_<hostname>`) to make sure that the next recognized (and consumed) line will be the Sun Microsystems copyright notice. If some other line occurs, it will be captured by the default report only rule (not shown) that fires when the `Reboot_in_Progress_...` context for the host is set. It also generates events that trigger rules to search for the startup messages from xntpd, inetd and sshd.

The `PairWithWindows` rules to detect daemon startup (e.g. xntpd, sshd, inetd) are not shown, but they consume one sequence of normal startup messages for 10 minutes after the reboot starts. They alert if the startup messages for these services are not found.

Rule 2 matches the Sun copyright notice if the prior event from the host was the genunix version declaration (matched by rule 1). It then creates a context that enabled the rule looking for the third line in the sequence.

Rule 3 matches the eri network interface detection event (again using a three trigger pair rule) and makes sure that is successfully configured at 100 Mbs within 10 seconds. It then sets a context to allow detection of the device lines that occur after the eri interfaces are detected.

I do not care about the order of the lines occurring after the network interface is detected, I just want to see that each device has been detected, so I used a coincidence detector in rule 4. Each rule that detects one of these devices generates a `CHECK ERI LINES <hostname>` event.

[2]Greatly truncated for publication. A more complete example is in the downloadable ruleset.

```
# Rule 1: detect reboot start and limit overall length of reboot
type= pair
desc= recognize start of reboot
...
pattern= ([^ ]*) genunix: .*SunOS Release 5.8 Version
action=add Reboot_in_Progress_$1 $0; create look_for_line_2_$1; \
        create reboot_failed 300 (write - reboot of $1 failed.); \
event xntp: reboot startup; event inetd reboot startup;\
        event sshd reboot startup
pattern2= $1 genunix: .* dump on /swapfile size
action2 = delete reboot_failed;
window=300

# Rule 2: recognize second line in reboot sequence
type= single
desc= recognize second line of reboot
...
pattern= ([^ ]*) genunix: .* Copyright 1983-2001 Sun Microsystems, Inc.
context = Reboot_in_Progress_$1 && look_for_line_2_$1
action= delete look_for_line_2_$1; create look_for_line_3_$1

# Rule 3: look for network interface and it must be up in <10 sec.
type=pair
desc = detect proper response of network card.
window=10
context=Reboot_in_Progress_$1 && look_for_line_32_$1
pattern= ([^ ]) genunix: .*eri([0-9]*) is /pci@1f,0/network
action= create failed_eri$2_config_$1 10 \
          (write - Failed to detect eri0 operation within 10 seconds); \
        delete look_for_line_32_$1; \
        create look_for_line_following_eri_detection_$1
ptype2=regexp
pattern2 = $1 eri: .* SUNW,eri$2 : 100 Mbps full duplex link up
action2 = create look_for_line_39_%1; delete failed_eri%2_config_%1; \
          add Reboot_in_Progress_$1 $0

# Rule 4: Check for other device recognition messages after
#   the eri card has been detected.
type=single
desc = all three devices following eri detection present
...
pattern= CHECK ERI LINES ([^]*)
context= found_ds0_$1 && found_uata0_$1 && found_devinfo0_$1
action = delete found_ds0_$1;delete found_uata0_$1; delete found_devinfo0_$1
```

Figure 21: A partial ruleset for analyzing a boot sequence.

4 Strategies to improve performance

One major issue with real-time analysis and notification is the load imposed on the system by the analysis tool and the rate of event processing. The rules can be restructured to reduce the computational load. In other cases the rule analysis load can be distributed across multiple systems or across multiple processes to reduce the load on the system or improve event throughput for particular event streams.

The example rule set from UMB utilizes a number of performance enhancing techniques. Originally these techniques were implemented in a locally modified version of SEC. As of SEC version 2.2.4, the last of the performance improvements has been implemented in the core code.

4.1 Rule construction

For SEC, construction of the rules file(s) plays a large role in improving performance. In SEC, the majority of computation time is occupied with recognizing events using Perl regular expressions. Optimizing these regular expressions to reduce the amount of time needed to apply them improves performance.

However, understanding that SEC applies each rule sequentially allows the reader to put the most often matched rules first in the sequence. Putting the most frequently used rules first reduces the search time needed to find an applicable rule. Sending a USR1 signal to SEC causes it to dump its internal state showing all active contexts, current buffers, and other information including the number of times each rule has been matched. This information is very useful in efficiently restructuring a ruleset

Using rule segmentation to reduce the number of rules that must be scanned before a match is found proves the biggest gains for the least amount of work.

4.1.1 Rule Segmentation

In August 2003, I developed a method of using SEC's multiple configuration file mechanism to prune the number of rules that SEC would have to test before finding a matching rule.

```
type=suppress
continue=dontcont
ptype=NRegExp
pattern=^[ABCD]
desc=guard for abcd rules

type=single
continue=dontcont
ptype=TValue
pattern=true
desc=guard for events handled by other \
     ruleset files
action=logonly
context = [handled]

type=single
continue=takenext
ptype=TValue
pattern=true
desc=report handled
action=create handled

<rules here>

type=single
ptype=TValue
pattern=true
desc=Guess we didn't handle this event after all
action=delete handled
```

Figure 22: A sample rule set to allow events to be filtered and prevented from matching other rules in the file.

This mechanism provides a limited branching facility within SEC's ruleset. A single criteria filtering rule is shown in Figure 22:

This rule depends on the `Nregexp` pattern type. This causes the rule to match if the pattern *does not* match. The pattern is crafted to filter out events that can not possibly be acted upon by the other rules in the file. In this example I show another guard that is used to prevent this ruleset from considering the event if it has been handled. It consists of a rule that matches all events[3] and fires if the `handle`

[3]The `TValue` ptype is only available in SEC 2.2.5 and newer. Before that use regexp with a pattern of `?`.

context is set. If it does not eliminate the event from consideration, I set the `handled` context to prevent other rulesets from processing the event and pass the event to the ruleset. If the final rule triggers, then the event was not handled by any rule in the ruleset. The final rule deletes the `handled` context so that the following rulesets will have a chance to analyze the event.

Note that this last rule is repeated in the final rule file to be applied. There it resets the `handled` context so that the next event will be properly processed by the rulesets.

In addition to a single regexp, multiple patterns can be applied and if any of them select the event, the event will be passed through the rest of the rules in the file. A rule chain to accept an event based on multiple patterns is shown in Figure 23. The multiple filter criteria can be set up to accept/reject the event using complex boolean expressions so that the event must match some patterns, but not other patterns.

The segmentation method can be arbitrary, however it is be most beneficial to group rules by some common thread such as the generator, using a a file/ruleset for analyzing sshd events and another one for xntp events. Another segmentation may be by host type. So hosts with similar hardware are analyzed by the same rules. Hostname is another good segmentation property for rules that are applicable to only one host.

The segmentation can be made more efficient by grouping the input using SEC's ability to monitor multiple files. When SEC monitors multiple files, each file can have a context associated with it. While processing a line from the file, the context is set. For example, reading a line from /var/adm/messages may set the `adm_messages` context, while reading a line from /var/log/syslog would set the `log_syslog` context and clear the `adm_messages` context. This allows segmentation of rules by source file. Offloading the work of grouping to an external application such as syslog-ng provides the ability to group the events not only by facility and level as in classic syslog, but also by other parameters including host name, program, or by a matching regular expression. Since syslog-ng operates on the components of a syslog message rather than the entire message, it is expected to

```
type= single
desc= Accept event if match2 is seen.
continue= takenext
ptype= regexp
pattern= match2
action= create accept_rule

type= single
desc= Accept event if match3 is seen.
continue= takenext
ptype= regexp
pattern= match3
action= create accept_rule

type= single
desc= Skipping ruleset because neither \
      match2 or match3 were seen.
ptype= TValue
pattern= true
context=  ! accept_rule
action= logonly

type= single
desc= Cleaning up accept_rule context \
      since it has served its purpose.
continue=takenext
ptype= TValue
pattern= true
context= accept_rule
action= delete accept_rule; logonly

<other rules here>
```

Figure 23: A ruleset to filter the input event against multiple criteria. The words "match2" or "match3" must be seen in the input event to be processed by the other rules.

be more efficient in segmenting the events than SEC.

Restructuring the rules for a single SEC process using a simple 5 file segmentation based on the first letter of the event using an 1800 rule ruleset increased throughput by a factor of 3. On a fully optimized ruleset of 50 example rules, running on a SunBlade 150 (128MB of memory, 650Mhz), I have seen rates exceeding 300 lines/sec with less than 40% processor utilization. In tests run under the Cygwin environment on Microsoft windows 2000, 40 rules produced a throughput of 115 log entries per second. This single file path of 40 rules is roughly equivalent to a segmented ruleset of 17 files with 20 rules each for a total of 340 rules, with events equally distributed across the rulesets.

Note that these throughput numbers depend on the event distribution, the length of the events etc. Your mileage may vary.

4.2 Parallelization of rule processing

In addition to optimizing the rules, multiple SEC processes can be run, feeding their composite events to a parent SEC. SEC can watch multiple input streams. It merges all these streams into a single stream for analysis. This merging can interfere with recognition of multi-line events as well as acting to increase the size of an event queue, slowing down the effective throughput rate of a single event stream. Running a child SEC process on an event stream allows faster response to that stream.

SEC's `spawn` action creates a process and creates an event from every line emitted by the child process. The events from these child processes are placed on the front of the event queue for faster processing.

These features allow the creation of a hierarchy of SEC processes to process multiple rules files. This reduces the burden on the parent SEC process by distributing the total number of rules across different processes. In addition, it simplifies the creation of rules when multi-line events must be considered, by preventing the events from being distorted by the injection of other events in the middle of the multi-line event.

SEC is not threaded, so use of concurrent processes is the way to make SEC utilize multiprocessor systems. However, even on uniprocessor systems, it seems to provide better throughput by reducing the mean number of rules that SEC has to try before finding a match.

4.3 Distribution across nodes

SEC has no built-in mechanism for distributing or receiving events with other hosts. However, one can be crafted using the ideas from the last two sections. Although this has not been tested, it is expected to provide a significant performance improvement.

The basic idea is to have the parent SEC process use ssh to spawn child SEC processes on different nodes. These nodes have rules files that handle a portion of the event stream. The logging mechanisms are set up to split the event streams to the nodes so that each node has to work on only a portion of the event stream. Even if the logs are not split across nodes, the reduced number of rules on each node is expected to allow greater throughput.

This can be used in a cluster to allow each host to process its own event streams and report composite events to the parent SEC process for cross-machine correlation operations.

4.4 Redundancy

As with distribution, SEC has no built in redundancy mechanism, however it should be possible to provide a measure of redundancy by duplicating rules files on a few hosts that receive the same event data. I have not tested this in its entirety, but the basic blocks have been tested.

Redundant SEC processes can generate heartbeat events for each other. SEC is able to detect when the events are missing and enable all of its rulesets.

While I have not set up the heartbeat mechanism, I have a ruleset that can be easily modified for that purpose. This ruleset was implemented because of a limitation in classic syslog. If a host is a syslog master and receives syslog events from child hosts, it cannot send its own events to another master syslog host without forwarding all of its child events. As a result each master loghost must correlate its own events.

Newer versions of syslog (e.g. syslog-ng) address this problem.

The ruleset suppresses all syslog events that do not originate on the syslog master host using a mechanism similar to the rule segmentation discussed in section 4.1.1. It requires manual intervention to turn on, but this could be modified to work in a redundant environment.

5 Limitations

Like any tool, SEC is not without its limitations. The serial nature of applying SEC's rules limits its throughput. Some form of tree-structured mechanism for specifying the rules would allow faster application. One idea that struck me as interesting is the use of ripple-down rulesets for event correlation[Clark2000] that could simply the creation and maintenance or rulesets as well as speed up execution of complex correlation operations.

As can be seen above, a number of idioms consist of mating a single rule to a more complex correlation rule to receive the desired result. This makes it easy to get lost in the interactions of more complex rulesets. I think more research into commonly used idioms, and the generation of new correlation operations to support these idioms will improve the readability and maintainability of the correlation rules.

The power provided by the use of Perl regular expressions is tempered by the inability to treat the event as a series of fields rather than a single entity. For example, I would prefer to parse the event line into a series of named fields, and use the presence, absence and content of those fields to make the decisions on what rules were executed. I think it would be more efficient and less error prone to come up with a standard form for the event messages and allow SEC to tie pattern matches to particular elements of the event rather then match the entire event. However, implementation of the mechanism may have to wait for the "One True Standard for Event Reporting", and I do not believe I will live long enough to see that become a reality.

The choice of Perl as an implementation language is a major plus because it is a more widely known language than C among the audience for the SEC tool this increases the pool of contributers to the application. Also, Perl allows much more rapid development than C. However, using an interpreted language (even one turned into highly optimized bytecode) does cause a slowdown in execution speed compared to native executable.

SEC does not magically parse timestamps. Its timing is based on the arrival time of the event. This can be a problem in a large network if the travel time cannot be neglected in the event correlation operations.

6 Other Applications

Although I have explored its use as a system/network administrator tool under Unix, SEC has other applications. One idea is to use SEC as a framework for software performance testing. However, it may not have sufficient granularity for jobs that require sub second timing. SEC has also been proposed as a mechanism for anonymizing log files [Brown5_2004].

Because SEC is written in Perl, it works on any Unix system as well as Microsoft Windows. I have tested it under the Cygwin environment. Although this paper has dealt with its use on Unix systems, it has been used to monitor the event log of a windows 2000 workstation and provide real-time notification of events using the windows messenger service. Using a tool such as snare [Snare], the windows event log can be turned into syslog events and correlated either on the Unix system, or on a Windows system by running a syslog daemon under windows (e.g. syslog-ng compiled under the Cygwin environment) to capture the event log and feed it to SEC.

7 Future Directions

Refinement of the available rule primitives and actions (e.g. the expire action) is an area for investigation. A number of idioms presented above are more difficult to use than I would like. In some cases these idioms could be made easier by adding new correlation types to the language. In other cases a mechanism for storing and retrieving redundant informa-

tion (such as regular expressions and timing periods) will simplify the idioms. This may be external using a preprocessor such as filepp or m4, or may be an internal mechanism.

Even though SEC development is ongoing, not every idea needs to be implemented in the core. Using available Perl modules and custom libraries it is possible to create functions and routines to enhance the available functionality without making changes to the SEC core. Developing libraries of add-on routines – as well as standard ways of loading and accessing these routines is an ongoing project. This form of extension permits experimentation without bloating SEC's core.

I would like to see some work done in formalizing the concept of rule segmentation and improving the ability to branch within the rule sets to decrease the time spent searching for applicable rules.

8 Availability

SEC is available from http://kodu.neti.ee/ risto/sec/

In addition to the resources at the primary SEC site above, a very good tutorial has been written by Jim Brown[Brown2003] and is available at: http://sixshooter.v6.thrupoint.net/SEC-examples/article.html

An annotated collection of rules files is available from http://www.cs.umb.edu/ rouilj/sec/sec_rules-1.0.tgz. This expands on the rules covered in this talk and provides the tools for the performance testing as well as a sample sshrc file for the ssh correlation example.

9 Conclusion

SEC is a very flexible tool that allows many complex correlations to be specified. Many of these complex correlations can be used to model[Prewett] normal and abnormal sequences of events. Precise modeling of events reduces both the false positive and false negative rates easing the burden on system administrators.

The increased accuracy of the model provided by SEC results in faster recognition of problems leading to reduced downtime, less stress and higher more consistent service levels.

This paper has just scratched the surface of SEC's capabilities. Refinements in rule idioms and linkage of SEC to databases are just a few of the future directions for this tool. Just as prior log analysis applications such as **logsurfer** influenced the design and capabilities of SEC, I believe SEC will serve to foster research and push the envelope of current log analysis and event correlation.

10 Author Biography

John Rouillard is a system administrator whose first Unix experience was on a PDP-11/44 running BSD Unix in 1978. He graduated with a B.S. in Physics from the University of Massachusetts at Boston in 1990. He specializes in automation of sysadmin tasks and as a result is always looking for his next challenging position.

In addition to his system administration job he is also an emergency medical technician. Over the past few years, when not working on an ambulance, he has worked as a planetarium operator, and built test beds for domestic hot water solar heating systems. He has been a member of IEEE since 1987 and can be reached at rouilj@ieee.org.

11 References

2swatch Homepage: ftp://ftp.sdsc.edu/pub/security/PICS/2swatch/README

Brown2003 Brown, Jim, "Working with SEC - the Simple Event Correlator", Web publication, November 23, 2003. URL: http://sixshooter.v6.thrupoint.net/SEC-examples/article.html

Brown5_2004 Brown, Jim, "SEC Logfuscator Project Announcement", simple-evcorr-users mailing list, 3 May 2004. URL: http://sourceforge.net/mailarchive/forum.php?thread_id=2712448&forum_id=2877

Clark2000 Clark, Veronica, "To Maintain an Alarm Correlator", Bachelor's thesis, The University of New South Wales, 2000. URL: http://www.hermes.net.au/pvb/thesis/

Finke2002 Finke, John, "Process Monitor: Detecting Events That Didn't Happen", USENIX Systems Administration (LISA 16) Conference Proceedings, pp. 145-154, USENIX Association, 2002.

Hansen1993 Hansen, Stephen E. and E. Todd Atkins, "Automated System Monitoring and Notification with Swatch", USENIX Systems Administration (LISA VII) Conference Proceedings, pp. 145-156, USENIX Association, November 1993. URL http://www.usenix.org/publications/library/proceedings/lisa93/hansen.html

logsurfer Homepage: http://www.cert.dfn.de/eng/logsurf/

LoGS Prewett, James E., "Listening to Your Cluster with LoGS", The 5th LCI International Conference on Linux Clusters: The HPC Revolution 2004, Linux Cluster Institute, May 2004. URL: http://www.linuxclustersinstitute.org/Linux-HPC-Revolution/Archive/PDF04/05-Prewett_J.pdf

logwatch Bauer, Kirk. Homepage: http://www.logwatch.org/

logsurfer+ Homepage: http://www.crypt.gen.nz/logsurfer/

NNM "Managing Your Network with HP OpenView Network Node Manager", Hewlett-Packard Company, January 2003. Part number J5323-90000.

Prewett Prewett, James E. via private email, March 2004.

ruleCore Homepage: http://www.rulecore.com

Sah02 Sah, Adam, "A New Architecture for Managing Enterprise Log Data", USENIX Systems Administration (LISA XVI) Conference Proceedings, pp. 121-132, USENIX Association, November 2002.

SEC Vaarandi, Risto. Homepage: http://kodu.neti.ee/ risto/sec/

SECman Simple Event Correlator (SEC) manpage URL: http://kodu.neti.ee/ risto/sec/sec.pl.html

SLAPS-2 Homepage: SLAPS-2 http://www.openchannelfoundation.org/projects/SLAPS-2

SHARP Bing, Matt and Carl Erickson, "Extending UNIX System Logging with SHARP", USENIX Systems Administration (LISA XIV) Conference Proceedings, pp. 101-108, USENIX Association, December 2000. URL: http://www.usenix.org/publications/library/proceedings/lisa2000/full_papers/bing/bing_html/index.htm

Snare InterSect Alliance. Homepage: http://www.intersectalliance.com/projects/SnareWindows/index.html

swatch Atkins, Todd. Homepage: http://swatch.sourceforge.net/

Takada02 Takada, Tetsuji and Hideki Koike, "MieLog A Highly Interactive Visual Log Browser Using Information Visualization and Statistical Analysis", USENIX Systems Administration (LISA XVI) Conference Proceedings, pp. 133-144, USENIX Association, November 2002. URL: http://www.usenix.org/events/lisa02/tech/takada.html

Vaarandi7_2003 Vaarandi, Risto, "Re: is this possible with SEC", simple-evcorr-users mailing list, 4 Jul 2003. URL: tt http://sourceforge.net/mailarchive/forum.php?thread_id=2712448&forum_id=2877