AccelOps Security Information and Event Management (AccelOps SIEM)

FAQs: AccelOps SIEM vs. other SIEM Solutions

1

I would like to have a single solution for both real-time log analysis as well as long term Log Management. My current SIEM product does not provide this capability — does AccelOps support this requirement?

Our optimized file-based event database coupled with parallel data management and analysis enables AccelOps customers to have a single solution for analyzing both real-time data and historical data. Computing and storage can be incrementally added without service disruption. In contrast, most SIEM vendors must purge and archive long term data to avoid overwhelming their real-time relation databases, necessitating the use separate tools — one set of tools to manage real-time, and another to manage historical data. This approach also has limitations for the amount of data that can be analyzed — so while the data stored may meet retention requirements, the ability to actually analyze and cross-correlate across the stored data is often severely limited.

2

I need to collect events from 100's of windows servers without agents and support the latest Windows version? Can AccelOps do that?

AccelOps can be deployed in clustered mode to accomplish this. The solution consists of many worker nodes and one supervisor node. The job of pulling windows logs from many servers via WMI is load balanced among the worker nodes. Each worker node is multi-threaded and can pull from many servers simultaneously. The events are parsed and indexed by each worker node and the correlation to trigger rules is done by the supervisor and worker nodes in a collaborative fashion. Additional worker nodes can be deployed if more windows servers need to be monitored and the system is running out of capacity.

3

My current SIEM system can only correlate and alert within 1 system. As I deploy event collection on servers or collect Netflow from routers, I need to deploy mulitple SIEM systems and need to correlate across them. Does your architecture support this?

The clustered mode AccelOps solution can do real time global cross-correlation across multiple supervisor and worker nodes. One simple way to do this would be to filter and forward all events to the supervisor node, but that would slow down the supervisor node. AccelOps employs a novel summarized information exchange mechanism where the worker nodes do the pre-processing of the events and only sends summarized values to the supervisor node, which can then do final analysis and trigger alerts. The entire category of rules can be parallelized this way by the AccelOps clustered system and provides customers a way to scale event processing and alerting

4

As I see an IP address in my dashboard or alert, I would like to know the user behind that IP and the user's network location if it is an internal one, or learn about the owner, domain etc from internet sites for external IPs. Can AccelOps do this?

AccelOps provides full identity and location information for IP addresses — both external and internal, and in real time. For internal IP addresses, AccelOps derives the identity information by combining Active Directory discovery, domain logon information, DHCP events, Wireless and VPN logons and the location information from Wireless and VPN logons and AccelOps own layer 2 discoveries. The challenge here is that each source of information is partial, e.g. DHCP address assignments provides (IP Address, MAC address and Host name), domain logon provide (IP Address, Host name, User name) etc. The various pieces need to be stitched together into one consistent identity and location entry and it should also dynamically reflect the changes as they occur as the user moves around. AccelOps has a novel in-memory database based approach for merging the pieces various identity and location information on a first and last seen time basis. This contextual information is available to the user for every IP address displayed on the user interface. AccelOps binds the user identity and location information to events to allow for historical analysis.

For an external IP address, AccelOps provides information such as geo-location, whois lookup and trace-route information. With a single click, administrator can find out whether this IP is a part of known spam databases using tools like SAN StormCenter, Cisco Senderbase or HoneyPot database.

5

How do I prioritize my alerts in AccelOps?

AccelOps has the notion of a business service that is a smart container of network devices, servers and applications serving a common business purpose. Every incident is tagged with the affected business service and can be used to prioritize incidents. AccelOps goes beyond traditional event severity by providing users business impact context to incidents.

6

As I investigate my incidents, I would also like to add additional context. For example, when there are lots of denied connections to a server, I would like to know the CPU and memory usage on the server as well as the changes that happened on the server in a preceding time period. Can I quickly do that in AccelOps UI itself or I have to jump to another console?

This is easily possible in AccelOps since all aspects of a device are monitored. All the user needs to do is to discover the device and set up monitoring. Then the basic system level CPU/memory/disk space/disk I/O/network interface utilization, the top applications consuming most resources on that server as well as the changes made on that server are all available within 1 click. The information is also kept up to date on a periodic basis.

7

I would like to customize various dashboards and create reports in various formats that I can show to my management. What capabilities does AccelOps provide for me to be able to do this?

AccelOps supports the exporting of real-time, historic and saved search results / reports in both PDF and CSV formats. The PDF reports contains multiple colored trend charts and can be customized with customer logos and custom notes. The CSV format can be exported directly into spreadsheet products or can be used to feed to other applications easily. Furthermore, saved reports are available as templates that can be used as dashboard widgets — enabling fully customized dashboards.

8

How flexible is your reporting framework? How many system reports do you ship with the product? Can I issue a simple Google-like search to find keywords in order to perform root cause analysis?

AccelOps features an advanced SQL-like search and cross-correlation engine with multiple patterns and advanced filtering and aggregation capabilities that can be computed in a distributed manner. This enables support of IT infrastructure, availability, performance, change and security scenarios, as well as allowing compliance requirements to be handled in a unified manner.

AccelOps ships with more than 1000 (and growing) of built-in and extensible reports spanning availability, performance, security and change management, as well as compliance and inventory.

We support simple keyword Google-like searches with operators such as AND, NOT, etc, and also feature the capability of searching through real-time event data using either structured (condition-based), or simple keyword queries

9

What is your database architecture and how scalable is it? Do I need to worry about purging and other issues? Can I add more storage capacity to the system as I expand my data centers? How good is the system performance when you have millions of events coming in and are performing queries on the data simultaneously?

AccelOps uses a hybrid database, storing events in indexed flat-files, and storing device configuration in an embedded commercial relational database (PostgreSQL). AccelOps has a patent-pending, multi-tiered, clustered architecture, where computing and storage can be seamlessly added to the cluster to increase performance and event storage capacity. This combination of proprietary database and parallel processing gives AccelOps the dual advantage of unlimited low cost storage and high event analysis performance that other monitoring solutions strive for.

10

How good is AccelOp's support for 3rd party devices? What's your framework for supporting 3rd party devices/applications? I want to add support for a new device — what do I need to do?

For our data center and cloud monitoring solution to be useful for availability, performance, security, change and compliance, we need to support the heterogeneous environment that reflects a real data center with best-of-breed equipment from various vendors. So we are committed to supporting third party software and device in a timely fashion.

Since AccelOps monitors all aspects of a data center, our support for a single device or application tends to be comprehensive, ranging from auto discovery to categorization and normalization of SNMP traps, syslogs, netflow, WMI metrics and other event/protocol formats concerning availability, performance, security and change. While a significant majority of Tier 1 and 2 vendors are already supported, we are continuously adding new device support and keeping the existing device support up to date.

There are some technical innovations that we have developed to accelerate the device support process. Typically, there are two ways to add device support - custom coding and scripts.

Custom coding involves parsing the device information within the shipping product code (Java or C++, within the main code or via an SDK in an agent). Scripts typically reside outside the main product code. The main tradeoff between custom coding and scripts is performance and flexibility. Scripts are flexible but custom coding gives you performance — a perl based program designed to parse Netflow data or firewall logs would certainly not be able to keep with the event rate for high-end routers/firewalls.

However, AccelOps has developed a unique XML based scripting language through which comprehensive device support can be added without sacrificing performance.

While XML based parsing definitions exist in other products (such as Splunk), AccelOps XML parsing language has the power of programming languages (e.g. if-then-else, switch-case, temporary variables, etc) that makes comprehensive device support possible. In addition, we have developed a XML compiler and execution environment that enables AccelOps the means to execute the XML code without losing performance. In fact all our device support is written using the parsing XML language.

- AccelOps device support library includes a large number (over 1000 and growing) of parsed event attributes that encompasses events and logs from various IT management domains. This enables flexible support for a wide range of devices. More importantly, this is done without losing event processing performance and storage efficiency.

These technological innovations enable rapid and flexible device support. All it takes is to modify an existing parser XML file or create a new parser XML file and add it to the AccelOps system. In this way, new versions of supported devices can be easily added since they simply often add a few new logs. AccelOps has a dedicated team focused on device support, allowing us to provide high-quality and timely coverage. The user community can also be easily leveraged — if a partner has introduced a parser, that parser's XML file can be redistributed to other customers

11

My current SIEM solution has very slow reporting in general, especially during high event rate processing. How is AccelOps' reporting performance? And what if I need even faster processing?

The slowness in query response times in many SIEM products often comes from the use of a relational database. While relational database are easy to build a SIEM system around, the read-only monitoring data is ill-suited for relational database because of the following reasons

- If data is inserted at high event rates (e.g. when dealing with firewall, netflow or Active Directory data), the database limit is quickly reached, causing the vendors to archive old data. In many systems, only a few months of data can be kept in a relational database. The effect is that another system needs to be brought up to analyze at the old data.
- Event data may require many attributes (over a few hundred) to be parsed; a relational database table with so many columns can be unwieldy and causes performance and storage inefficiencies (also known as degradation and bloat).
- Parallelizing a relational database for faster query performance is a non-trivial matter, both in terms of cost and implementation and maintenance complexity

On the other hand, the discovered information about devices, systems and applications (so called CMDB) is highly structured, updated often, data that merits a relational database.

AccelOps has developed a hybrid data management system that stores unstructured event data in flat file based database and structured CMDB data in an embedded commercial relational database (PostgreSQL). A data management layer unifies the two data management technologies and presents a single relational database like interface and the best of both worlds is achieved. As an embedded RDBMS, the system does not require administrative tuning / index optimization.

More importantly, the ability to store events in a flat file database also enables query parallelization and solves the slow reporting problem. In clustered mode, AccelOps solution is deployed in a hierarchical supervisor-worker setup as explained previously. The supervisor node divides a query into many sub-queries, distributes the sub-queries to the worker nodes, and creates the final query result by combining the results from the various worker nodes. Since the flat files are stored in NFS on a separate system, instant query response time reduction can be obtained by simply bringing up additional worker nodes.

12

I need a flexible rules architecture so that I can change firing frequency, and create more sophisticated rules to catch security incidents. For example, "3 login failures followed by a success within a 10 minute time window", or "multiple login failures not followed by a success to the same system within a 1 day time window". Does the AccelOps rule architecture support this?

AccelOps contains a sophisticated rule framework that can support anything from simple threshold performance rules, to highly complex security rules, all with a simple user interface. It supports the following constructs:

- More than 1000 event attributes with which to form rule conditions
- Operators such as equals, greater than, IN, CONTAINS, BETWEEN, IS and their negative conditions
- Ability to create multiple sub-patterns then combine them using the temporal operators: AND, OR, FOLLOWED_BY, OR_NOT, AND_NOT, and NOT_FOLLOWED_BY
- Ability to create exceptions to rules in order to fine-tune their output
- Ability to exclude rules from firing during specific time ranges
- Ability to send business policy based alerts on resulting incidents, thus sending the right kind of alert to the right person. Alerts can be sent via email, SMS, SNMP Traps or XML via HTTP

AccelOps supports both simple and advanced work-flows when creating or editing rules.

accelops

Web: www.accelops.com
Email: info@accelops.com
Tel: +1 (408) 970-9668
Fax: +1 (408) 970-9666

FREE TRIAL DOWNLOAD
www.accelops.com/download