

# Network Computing

MAY 25, 2006 | WWW.NWC.COM

*For IT By IT*

[ Product Analysis ]

# SIM

WE POPULATED OUR CHICAGO NEOHAPSIS LAB WITH EIGHT SECURITY INFORMATION MANAGEMENT PLATFORMS. AN INDUSTRY VETERAN TOOK THE CROWN, BUT NEW RIVALS ARE ON THE PROWL **BY GREG SHIPLEY**



**Security teams ask a lot of today's SIM platforms.** We want them to take data from a dizzying array of sources—ArcSight ESM alone supports more than 200 log formats—then determine what's important out of the collected information, store terabytes of data online for analysis and trending, all while meeting both real-time analysis and historical reporting needs. Complicating matters, not all organizations ask for the same things; some are focused on compliance and reporting, while others are more concerned about monitoring and incident handling. Any one of these presents a difficult challenge, yet today's security information management platform vendors seem interested in tackling most, if not all, scenarios.

Those familiar with the SIM arena will note that our review has a number of new players, while a few old faces are conspicuously absent. We invited 13 vendors to send their latest offerings to our Chicago Neohapsis Real-World Labs®. We required that products support at least 12 log formats natively, centralize log data and provide an analysis component for reducing the amount of information a human operator must process. ArcSight, High Tower Software, LogLogic, Network Intelligence, OpenService, Q1 Labs, SenSage and Symantec accepted. Long-time players eSecurity (now owned by Novell), Intellitactics, netForensics and Guarded.Net (now owned by IBM) declined, as did Cisco Systems. CA indicated interest but never sent product, and we didn't learn of eIQnetworks and TriGeo Network Security until after our testing deadline.

Of the original pack of well-known SIM vendors, only ArcSight and Network Intelligence participated, which we take as a sign that the functionality gap between leaders and followers might be growing. We wonder if some older SIM vendors have fallen even further behind.

Our evaluation requirements focused on data transportation, storage, reporting, and both real-time and forensic-analysis tools. Our test environment consisted of varying types of devices, ranging from Windows 2000, 2003 and Linux systems to firewalls, VPN concentrators, IDSs, and infrastructure devices, such as routers and switches. Our plan was to bring each system into our lab, have vendor reps install and configure it, then incorporate the product into our day-to-day operational activities. We were hoping the process would be straightforward, but what ensued was a months-long engagement during which we learned about differing UIs, terminology, architectures, scripting languages, correlation strategies and investigative toolsets.

## A Landscape Revised

**One immediately obvious change:** We could deploy some of the SIM products in a single day, providing the infrastructure had been prepped—devices have logging enabled and configured, with logs already forwarded to a central destination. Three years ago, when we tested SIMs, a one- or even two-day installation would not have been possible. Yes, there's still a large difference between bringing several dozen devices into a logging infrastructure versus several hundred or thousand, but even getting a base product up, running, and receiving and processing logs in less than eight hours is a huge step forward. We were thrilled with the progress here.

Also worth noting is the number of vendors that have gone the turnkey route. Three years ago, the vast majority offered strictly software-delivery models. Today, most of our participants sent their products on systems or appliances, which helps customers avoid additional hardware, OS and database licensing costs while greatly reducing installation headaches. We understand that in large, multi-terabyte SAN-enabled deployments a software-based system might be desirable, but we question how many of those situations exist. From our perspective, as long as OS patching is handled in an efficient and timely manner, we like the turnkey model. It's cleaner. Only ArcSight and SenSage lack turnkey delivery products, so clearly we're not alone.

## Digging In

**This is the third time** we've built out a logging environment for a SIM bake-off, and the experience allowed us to avoid some pitfalls.

For starters, we knew that using a combination of *syslog-ng* with a UDP syslog relay would reduce the number of steps required for troubleshooting (see "How We Tested SIM Products," page 6). For the uninitiated, *syslog-ng* is an open-source package that was created by Balazs Scheidler as a more functional replacement to the original Unix *syslog* daemon. *syslog-ng* uses the *syslog* protocol, but it offers a healthy amount of additional functionality—for example, support for *syslog* over TCP—and has gained immense popularity, in part for its flexibility when routing and retransmitting *syslog* messages.

By implementing *syslog-ng* and sending all our log data to

REAL-WORLD  
LABS®

# REPORT CARD

## SIM Products

	ArcSight ESM	Network Intelligence enVision	High Tower Software Security Event Manager	Q1 Labs QRadar	Symantec Security Information Manager 9500	LogLogic ST 3000 and LX 2000	SenSage Enterprise Security Analytics	OpenService Security Threat Manager 3.5
<b>REAL-TIME FEATURES</b>								
User interface (20%)	4.5	3.5	4	3.5	3.5	3	3	2
Event correlation (15%)	5	4	3	4	3.5	1	3	2
<b>CORE FEATURES</b>								
Device support (10%)	5	4	3	3	3	2	4	4
Data management (5%)	4	4	3	3	2	4	3	3
Device provisioning (5%)	4	3	5	3	2	4	2	2
Transport capabilities (5%)	5	3	3	3	3	3.5	3	2
<b>PRICING (25%)</b>	3	4	4.5	4	3	4	2	3
<b>REPORTING (15%)</b>	4	3	2	3	3	3.5	4	3
<b>TOTAL SCORE (100%)</b>	<b>4.15</b>	<b>3.65</b>	<b>3.53</b>	<b>3.50</b>	<b>3.08</b>	<b>3.05</b>	<b>2.95</b>	<b>2.65</b>

A≥4.3, B≥3.5, C≥2.5, D≥1.5, F<1.5  
A-C GRADES INCLUDE + OR - IN THEIR  
RANGES. TOTAL SCORES AND WEIGHTED  
SCORES ARE BASED ON A SCALE OF 0-5.

B+

B-

B-

B-

C+

C+

C

C-

**TRANSPORT CAPABILITIES** rates how well the product moves data from point A to point B.

**DATA MANAGEMENT** reflects how well we could manage data once it's in the product; are there DBA-like UI tools or is everything manual?

Customize the results of this report card using the Interactive Report Card®, a Java applet, at [www.nwc.com](http://www.nwc.com).

a relay host we ensured all the SIM products received the exact same data, and verified that device data was being received. We didn't want to get stuck in any finger-pointing contests with vendors claiming their devices weren't receiving logs.

On the transportation front, we faced an all too familiar dilemma: Either place agents on all our log sources to move data off them securely, or go the classic, unreliable route of using syslog over UDP (see "Stuck on Syslog," page 4). In the case of appliances, such as Cisco PIX firewalls, we didn't have a choice—you can't run agents on a PIX. Grudgingly, we opted to go with syslog because we couldn't stomach the idea of trying to deploy agents everywhere. Had we gone the agent route, however, ArcSight ESM would have been our best

option because its collectors not only encrypt traffic, they also perform bandwidth throttling and batch transfers. These features come in handy when dealing with remote sites and limited WAN connectivity. In really large organizations, we could see the advantage of going agentless in some areas and using agents in others. We're skeptical of vendors that claim one model is superior; there's a time and place for each.

On the storage front, our first order of business was assessing our log volume sizes (see "The Windows Logging Headache," page 7). We doubted that we'd hit the multi-terabyte ceilings that plague many SIM deployments, but we wanted to be certain. Once we had all of our devices speaking syslog, we brought a syslog-ng Linux system online to serve as a temporary repository. After a few weeks of log col-

## SIM VENDORS AT A GLANCE

### PARTICIPATING COMPANIES

Company name	Year founded	Product name	Year launched	Key customers	News
ARCSIGHT (PRIVATE)	2000	ESM	2002	Defense Information Systems Agency, Capital Blue Cross, Iberdrola, Network Appliance, Unisys	Added support for Oracle Database Vault; added five global resellers
HIGH TOWER SOFTWARE (PRIVATE)	1999	Security Event Manager	2005	Undisclosed	Launched the 3200 series of appliances
LOGLOGIC (PRIVATE)	2003	LogLogic LX 2000; LogLogic ST 3000	2003	Yahoo, Royal Caribbean, Perot, Northwestern Memorial Hospital	Launched Project Lasso, an open-source project to monitor Windows events
NETWORK INTELLIGENCE (PRIVATE)	1996	enVision	2002	Undisclosed	Launched real-time data-mining application; hired new VP of global channel sales
OPENSERVICE (PRIVATE)	1994	Security Threat Manager 3.5	2002	VeriSign, Verizon Wireless, Security Bank of Kansas City, Young America	Ported its Security Management Center software to IBM BladeCenter
Q1 LABS (PRIVATE)	2001	QRadar	2001	Harvard University, Kerr-McGee, NASA, U.S. Army	Launched QRadar 3102 appliance
SENSAGE (PRIVATE)	2000	Enterprise Security Analytics	2001	BT Infonet, Credit Suisse, Duke University, Lehman Brothers	Announced integration with EMC Smarts for log management
SYMANTEC (SYMC)	1982	Security Information Manager 9500	2001	Undisclosed	Announced Q4 2005 \$118.8 million net profit, a slight decline from previous year

Source: Company reports, Yahoo.com

### NONPARTICIPATING COMPANIES

Company name	Year founded	Product name	Year launched	Reason for nonparticipation
CA (CA)	1976	eTrust Security Command Center	2003	Testing parameters
CISCO SYSTEMS (CSCO)	1964	Monitoring, Analysis and Response System (MARS)	Undisclosed	Undisclosed
EIQNETWORKS (PRIVATE)	2001	Enterprise Security Analyzer	2005	Learned of review too late
IBM (IBM)	1911	GuardedNet, Netcool/NeuSecure	Undisclosed	Timing: IBM acquisition of Micromuse
INTELLITACTICS (PRIVATE)	1996	Intellitactics Security Manager	2006	Timing: new release not available for testing
NETFORENSICS (PRIVATE)	1999	nFX Open Security	1999	Timing: new version not available for testing
NOVELL (NOVL)	1983	Sentinel	Undisclosed	Timing: Novell acquisition of eSecurity
TRIGEO NETWORK SECURITY (PRIVATE)	2001	TriGeo SIM	2002	Learned of review too late

Source: Company reports

lections we could estimate our average weekly log volume sizes. We had fewer than 40 devices generating log data (though some were quite chatty). They delivered between 40 and 60 events per second on average, about 3 GB of data a week.

Organizations should go through this exercise to understand how much data they'll have. We've spoken to organizations that generate more than 1 TB a month. Knowing how much data you'll have and how long to keep it, and understanding how much to keep online versus offline are critical to a successful SIM deployment. We knew going into this review that 500 GB to 1 TB of back-end storage would suit our needs just fine.

Also critical is an understanding of the limitations of back-end systems. Network Intelligence and SenSage have some advantages in this department, as they've moved away from conventional relational database technology to more proprietary warehousing mechanisms (see "Is RDBMS Bad in the SIM World?," page 10). The avoidance of expensive table reindexing and removal of unnecessary RDBMS features, such as record locking, could lead to much better performance when getting into multi-terabytes of data. Given our environment's size, we didn't have performance problems with any of the products we used.

Once we had the SIM products installed and data flowing freely, we turned our attention to operational needs. Our requirements revolved around two action areas: monitoring a select number of our critical assets and using investigation and query tools to determine if events ID'd by our monitoring efforts were real threats.

Monitoring capabilities are heavily dependent on correlation capabilities, especially when you're dealing with hundreds of events per second. If the SIM can't summarize and dismiss the vast majority of events coming into your console, you're fighting a losing battle. By analyzing information from firewalls, IDSs, authentication services and system hosts, good correlation rules can help identify what's of concern and what isn't. On the correlation front, ArcSight ESM's rule sets are the most powerful, but Q1 Labs QRadar's correlation logic is by far the easiest to use. Network Intelligence's enVision and Symantec's Security Information Manager 9500 also can perform a respectable amount of real-time correlation. By comparison, we found the prospect of authoring rules with SenSage's Enterprise Security Analytics falling somewhere between extremely painful and impossible; the product has a long way to go in the ease-of-use department.

For our tests we designed several custom correlation rules. Only a select number of systems and services in our lab are accessible from the outside world, for example, and those openings in our firewalls represent our primary attack surface; we keep a sharp eye on them at all times. One rule we devised was to cross-reference firewall-allow statements with successful authentication sessions. One potential sign of a successful

service attack would be an inbound connection that didn't result in a normal user session. Building this rule took a little time, but we succeeded with products from ArcSight, Q1 Labs and Symantec. We struggled with all of the others, and eventually threw in the towel for this particular rule.

Classifying individual systems also helps the prioritization effort. A system responsible for financial transactions or housing R&D data, for example, is more critical than the file server that holds marketing literature. When we went to classify our own systems we found High Tower Security Event Manager's asset-classification process easiest and most useful, but enVision and ArcSight ESM also offer asset-weighting functions. After setting up classification weightings, we're not sure how we ever survived without them.

Finally, when it comes to reporting features and general interface usability, most of the products need significant work. If you're spending any serious amount of time behind a console, you want to easily drill down, perform ad hoc queries on user names and IP addresses, actually use the pretty graphs by being able to click on them, and avoid frustration while navigating the UI. The products from High Tower and ArcSight are the two easiest products to use, interface-wise, with ArcSight

## STUCK ON SYSLOG

When we started reviewing SIM products back in 2002, one of the issues we found most alarming was the "lowest common denominator" used by most log-transportation efforts: syslog over UDP. It seemed almost comical to us that we, as a community, have invested millions in high-speed firewalls, advanced signature and inspection engines, and encryption solutions only to tunnel alert messages using a disturbingly low-tech, in-the-clear and unreliable delivery mechanism.

Here we are, four years later, and little has changed. Syslog over UDP is still the common alert-transport method for many devices. Sure, there are alternatives, such as Check Point Software's OPSEC and the technology behind utilities such as ArcSight's flex-agent, but none has taken hold as the transportation standard with which to work. We're still stuck with syslog.

We caught up with Chris Lovin, co-chair of the IETF syslog working group, to find out if there's hope. Fortunately, the syslog group has been busy; it has a number of initiatives set to become standards in 2007. For starters, the working group formally documented the syslog protocol. Implementations of syslog have evolved over the years, Lovin says, but different versions have done different things, and there hasn't been a formal standard.

One has now been written and is in the process of becoming finalized. In addition, the transport mechanisms (both TCP and UDP) have been officially documented, and a draft standard is in the works to transport syslog messages securely using TLS, which is arguably one of the most badly needed components in the messaging equation.

Assuming these initiatives make their way through the IETF standards process, it will then become our job to ensure they enter the real world. We must start demanding that our application, OS and device vendors adopt them. Here's hoping we won't be talking about this issue the next time we review SIMs.

ESM's UI far more flexible and comprehensive than rivals'. EnVision has made advances in the UI department, but it still has a way to go. The field goes downhill from there. With the exception of LogLogic's ad hoc querying ability, which made some tasks ridiculously easy, the rest of the SIMs tested need serious work in the UI department.

In the end, we awarded ArcSight ESM our Editor's Choice, followed closely by enVision, High Tower and QRadar. In the raw functionality department, ArcSight ESM is the most mature product on the market, but it gets expensive fast in large environments. High Tower SEM's simplicity is attractive—we'll watch it closely over the coming year. Symantec also has a unique offering—its integration with its DeepSight data feeds lets organizations use information and data from other global resources. No other product offers this cross-organization data-sharing capability.

One thing to note about as-tested pricing scores, which are based on the SIM collecting and analyzing log data from just fewer than 40 devices: Products with a per-device pricing model become exponentially more expensive in large organizations. Although our deployment was relatively small, quadrupling our device load with some of the appliances, like High Tower's or LogLogic's, wouldn't have cost us a dime more. However, the same level of expansion with the software from ArcSight, Sensage or Symantec would have cost a pretty penny. This is why we weighted price at 25 percent of the score—the issue isn't just upfront outlay, but scaling costs (as-tested pricing can be found in the features chart, page 9).

**ArcSight ESM ArcSight's SIM suite, marketed under the guise of ESM, is the most comprehensive SIM tool we tested, and it earns its second consecutive first-place finish (it took top honors in our 2003 review).**

ESM sports an extremely flexible “connector” framework capable of bandwidth throttling and batch queuing; both Web and Java console access; reporting and incident-analysis tools; support for a wide range of device types, including more than 120 products from 60-plus vendors; and a comprehensive real-time correlation capability that we found helpful. Although the suite is not without its faults, and could be pricey, ArcSight ESM remains a few steps ahead of the competition (though we do see Network Intelligence's and High Tower's products gaining ground).

Our ArcSight deployment ran on three Windows 2003 systems: One housed the “manager” and served as the destination for all correlation and inbound data, one hosted the Oracle database for our back-end storage, and one ran the Java console. The console was a bit of a pig—it didn't run smoothly until we gave it 2 GB of RAM!

ArcSight ESM comes bundled with prebuilt correlation rules that helped us make sense out of a wide range of otherwise confusing activity. Rules flagging potential worm outbreaks due to traffic patterns, identifying suspicious activities,

such as “multiple firewall denies followed by an allow,” and spotting failed login attempts across multiple systems are just a few among a wide prepackaged set.

We did, however, have some false-positive problems with the stock rule set. For example, the “worm outbreak” rule seemed to misfire frequently. The first few times, the alert sent us into a bit of a panic as we scoured our network for signs of automated nastiness, but we soon discovered that the rule was being triggered by some overzealous Web surfing—an end user's rapid connects to multiple destinations through the firewall confused this particular ArcSight ESM correlation rule. We also ran into problems with “dark space” alarms (traffic communicating to/from our network to IP ranges that are supposedly unallocated or “dark”) that were attributable to outdated IP lists. We tuned these rules and eliminated the false positives, but we found it ironic that our primary alert consolidator was creating more alerts.

ArcSight ESM also offers a wide set of preconfigured dashboards, as well as a dizzying array of tools to monitor and manage the system. There's an entire dashboard, for example, complete with real-time graphs and tables, for simply monitoring database health and activity. ArcSight also uses a taxonomy-based approach that puts event and alert data into a generic classification, or bucket. A PIX deny event will be placed in a general firewall rule violation bucket, for instance. Not only does this simplify the process of creating correlation rules, it has the neat side effect of letting organizations swap out device vendors—you can replace Check Point firewalls with those from Juniper without worrying about your correlation rules breaking.

Our complaints about ArcSight are restricted to two areas: pricing and ad hoc querying limitations. On the pricing front, we found ArcSight's model cumbersome. The company charges per CPU for its manager, per Java console, per connector type and per device, and has some additional add-ons, such as visu-

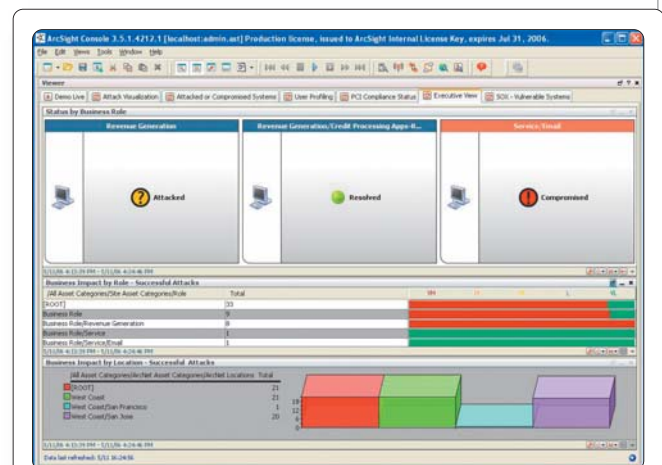


seemed to misfire frequently. The first few times, the alert sent us into a bit of a panic as we scoured our network for signs of automated nastiness, but we soon discovered that the rule was

being triggered by some overzealous Web surfing—an end user's rapid connects to multiple destinations through the firewall confused this particular ArcSight ESM correlation rule. We also ran into problems with “dark space” alarms (traffic communicating to/from our network to IP ranges that are supposedly unallocated or “dark”) that were attributable to outdated IP lists. We tuned these rules and eliminated the false positives, but we found it ironic that our primary alert consolidator was creating more alerts.

ArcSight ESM also offers a wide set of preconfigured dashboards, as well as a dizzying array of tools to monitor and manage the system. There's an entire dashboard, for example, complete with real-time graphs and tables, for simply monitoring database health and activity. ArcSight also uses a taxonomy-based approach that puts event and alert data into a generic classification, or bucket. A PIX deny event will be placed in a general firewall rule violation bucket, for instance. Not only does this simplify the process of creating correlation rules, it has the neat side effect of letting organizations swap out device vendors—you can replace Check Point firewalls with those from Juniper without worrying about your correlation rules breaking.

Our complaints about ArcSight are restricted to two areas: pricing and ad hoc querying limitations. On the pricing front, we found ArcSight's model cumbersome. The company charges per CPU for its manager, per Java console, per connector type and per device, and has some additional add-ons, such as visu-



The ArcSight ESM Javd UI was the best in the review. It requires 2 GB of RAM, but provides stunning real-time event views.

alization suites and ticketing integration, that cost—you guessed it—extra.

By comparison, companies such as High Tower charge per appliance, period. ArcSight has the luxury of being the leader in the space and has been able to get by with this pricing structure, but we wonder how tolerant the market will be moving forward with an increased level of competition.

Finally, on the querying front, we have some concerns with ArcSight ESM's implementation. To perform ad hoc queries on, say, an IP address, we had to configure a data feed and range, pull a large dataset from the database to the console, and then perform a search on that local dataset. ArcSight's console applet does a lot of the work, but the task still requires multiple steps, can be confusing, and is often slow and inefficient given how much data you're pulling across the wire. By comparison, Log-Logic's ad hoc querying abilities go directly to the database, are incredibly simple and quick.

### Network Intelligence enVision Network Intelligence's enVi-

**B-** sion, a Windows-based, turnkey SIM, has come a long way since we last tested it. The product now includes more advanced correlation capabilities, customizable dashboards, and a wider range of supported devices and transport protocols. One additional selling point that Network Intelligence promotes is that enVision doesn't use a RDBMS under the hood to manage data; rather, the company built its own storage technology, called LogSmart IPDB, that does away with much of the unnecessary overhead found in most commercial RDBMSs. Although we couldn't increase the data in our test environment to the terabyte-size range required to measure this feature, we believe the move away from conventional RDBMS could be good if done right.

We brought enVision online in a few hours. In less than half a day we had the product up, running and taking all the feeds from our data sources. However, the device provisioning process was not as smooth as it was with High Tower's prod-

uct; we had to define and configure every one of our log sources because the product lacks an autodetection mechanism. This wasn't a huge drawback for us, because we had fewer than 40 nodes sending data into the SIM environment, but organizations with hundreds or thousands of devices might be less tolerant of the provisioning process.

After spending months with the product, we believe enVision's selling point is that it covers most feature requirements reasonably well: It has a real-time console, can perform correlation functions, includes basic reporting components, and we could get around the product without too much headache. However, its feature set is not as comprehensive as that of ArcSight ESM, and its UI isn't nearly as polished.

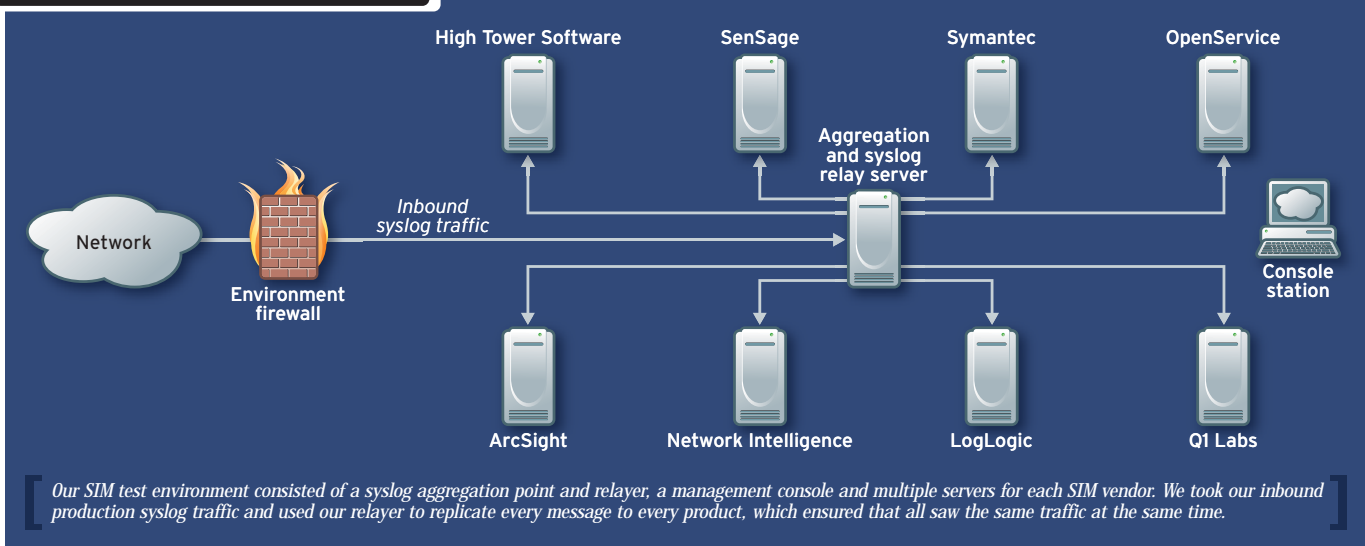
But given that enVision can get most jobs done and its no-hassle pricing model, we anticipate Network Intelligence will continue to give other SIM providers a run for their money.

### High Tower Software Security Event Manager This is

**B-** the first time we've used a High Tower appliance, and we're impressed. For starters, Security Event Manager (SEM) had the easiest deployment of any product we tested. Not only were configuration and setup a breeze (and finished in a morning), the device provisioning process was an absolute joy. Essentially, when we began pointing our syslog streams at the High Tower system, it intelligently identified each source, flagged it as a potential asset candidate, and made a reasonable guess as to what it was—for example, Cisco PIX or Linux host. The UI consolidated these device guesses under a "candidates" tab, and we simply approved or disapproved them as legitimate log sources. High Tower's well-thought-out approach saved us time and energy in our 30-some device environment—its impact in reducing deployment times in a several-hundred or several-thousand device environment would be significant.

SEM also had the most no-hassle interface of the products we tested. In fact, we found ourselves gravitating toward it

## How We Tested SIM Products





frequently. During testing we misconfigured a relay host, for example, resulting in the transmission of large amounts of data onto our outbound Internet links—a move that sent our latency through the roof. We immediately noticed that our events per second spiked from an average of 100 or so to more than 2,000, right from the SEM dashboard. Now, that's not to say that the other products can't be configured to chart event-per-second ratings, but because SEM has such a clean interface, it became our go-to product for many investigative efforts.

SEM's biggest shortcomings are on the reporting and real-time correlation fronts. Its reporting capabilities are simple at best. Canned reports are customizable, but what comes out is extremely basic. As for correlation, SEM is currently not in the same league as ArcSight ESM, enVision, QRadar and Symantec's Security Information Manager. SEM 3.1 (the latest shipping version) comes with 20 prebuilt correlation rules that are quite useful, but we couldn't build new rules from scratch. The only thing we could do is provide some tuning parameters for those 20. We saw a working demo of Version 3.2, due out later this year, and High Tower added the ability to customize real-time correlation rules, addressing our biggest concern with the product.

High Tower also has a very cut-and-dry pricing model; it ships turnkey systems that are sized based on two factors: event-per-second load and storage capacity. That's it. For organizations willing to sacrifice some flexibility and features

for ease of use and deployment, High Tower SEM will be hard to beat.

**Q1 Labs QRadar QRadar started off less as a SIM tool and more as a network traffic anomaly detection suite,**

**B**ut Q1 has added features, such as a correlation engine, the ability to take data feeds from multiple device types and a reporting engine, that make it competitive in the SIM market.

Q1 and OpenService are the only vendors in this review that support both software-only and turnkey delivery models. We opted to go the turnkey route, and Q1 delivered a 2102 QRadar server appliance and a 1101 flow collector appliance. We used the flow collector only to gain greater visibility into our environment; all the SIM-related smarts were in the 2102 server. The 2102 takes flow data from infrastructure devices, but gives customers the flexibility to deploy a standalone collector—a nice option and one that we used.

QRadar's correlation language is one of the easiest to learn. Unlike the rule-building sets in products from ArcSight and Symantec, QRadar's scripting language is probably as close to English as you're going to get. It's also somewhat hierarchical and recyclable, letting you use what Q1 calls "building blocks" to string logic sequences together. Although all the products required time to learn the correlation techniques, we were most comfortable using QRadar's methods.

We also liked that our logs could be viewed easily; we didn't have to rummage through the UI to get to them. This let us pinpoint a problem with QRadar's parsing of our PIX 7.0 logs, and Q1 turned around a fix in 24 hours. Had we not had easy access to the raw log data, it would have taken us a lot longer to pinpoint the problem. Most of the products we test-

## THE WINDOWS LOGGING HEADACHE

One challenge Microsoft Windows admins face is finding and sending Windows events to a central repository that isn't Microsoft-centric. Unfortunately, we still live in a syslog world (see "Stuck on Syslog," page 4), and though most OS and device vendors support at least syslog, Microsoft does not. Also challenging are third-party apps running on Windows systems that don't use the native Windows event-logging APIs. We discovered during our tests that our Vasco authentication system couldn't log events to Windows Event Log. Astounding.

Fortunately, a number of third-party commercial and open-source tools will forward native Windows event logs to syslog-based receiving points. If you can

get data into the native Windows event log, whether system, application or security, NTsyslog can move your data further down the path.

When we embarked on our SIM journey we opted to go the third-party route, and assumed that once we got our Windows events translated and forwarded via syslog things would go smoothly. Wrong. We soon discovered that once you enter the world of Windows event logs-to-syslog translation, there's no one format used. One Windows syslog forwarding mechanism will format data differently from the next, or will use different delimiters, wreaking all sorts of havoc.

After realizing this problem halfway through our deployment, we decided to

go with Snare, as it seemed to be the Windows events-to-syslog mechanism with which most SIM vendors were familiar. But even that road had a few bumps, as we discovered that Network Intelligence uses a different delimiter in Snare than does High Tower.

There are a few things we should be doing as a community to help solve these problems. First, we need to continue to pressure third-party application providers to give us options for logging to native Windows event logs. Second, we need to keep a sharp eye on the IETF efforts that standardize and secure syslog messages and transport methods. Last, once those standards are ratified, we need to pressure vendors—including Microsoft—to support them.



ed let organizations store raw events in addition to their normalized counterparts, but the raw data isn't always easily accessible.

Our biggest complaint with QRadar is with its UI; it's not a lot of fun to try to get things done within QRadar. We struggled to efficiently sift through poorly constructed HTML tables, and we often found ourselves unsure of which pane to use for what. By comparison, ArcSight ESM's Java console just crushes most HTML UIs in the usability department; Web UIs simply don't perform certain tasks well. We're not normally big fans of Java consoles because most of them are slow and kludgy, but given the choice between brutally painful HTML UIs and Java consoles that need 2 GBs of RAM, we'll opt for the Java—RAM is cheap. The console shortcoming won't be a big deal for those simply running some reports occasionally, but analysts spending hours in front of the console will likely find themselves getting uncomfortable over time without some improvements.

### Symantec Security Information Manager 9500



**Symantec entered the SIM space quietly when it acquired a small vendor named Mountain Wave back in 2002,** and only recently started marketing its SIM

offering aggressively as the revamped Symantec Security Information Manager 9500. Symantec shipped us a 9500 appliance, which we got up and running in a day, albeit with some minor struggles.

Unlike High Tower and LogLogic, Symantec is still bundling individual device collectors and parsers separately. So to bring a new class of device online you must first install the proper parsing code—a task that is manageable, but cumbersome. We found this approach particularly odd given that Symantec is using an appliance—why not just bundle everything on the device and be done with it? The 9500 wasn't as smooth as High Tower SEM on the device-provisioning front, either; we had to identify all our logging sources manually. Perhaps minor points given the full range of features, but if we were forking out \$60,000-plus for a SIM product we'd want these problems resolved.

Moving past installation nuances, the product is relatively mature in the areas of real-time analysis, correlation and reporting. The 9500 supports configurable dashboards, an incident-handling and ticketing system, and a basic set of canned reports. Like the products from ArcSight, Q1 and Network Intelligence, Symantec's offering also supports a GUI for creating custom correlation rules. The correlation rules editor and logic were a little awkward at first, until we understood more about the "AND and OR" modeling Symantec uses to build rules.

Symantec's pricing model is a combination of starting price, per-device price and annual support costs that include threat data feeds. The MSRP base is on par with that of offerings from companies such as High Tower, but Symantec also charges additionally per device, which makes the model less convoluted than ArcSight's, but more complicated than that of many of the other vendors.

Symantec does offer one unique feature: Through mining data sets from its DeepSight service, Symantec bundled an intelligence feed so that the 9500 can download and integrate known IP ranges and attackers into its internal data sets. This gives you some global perspective of top threats active on the Internet. For example, all the SIM products we tested could detect port sweeps and generic network scanning. However, by integrating the Symantec threat feed and incorporating it into live watch lists, the 9500 could identify common attackers on the Internet and give us additional information about frequency, timelines, when they were first spotted and more. We found this information useful when determining which probes to investigate further.

### LogLogic ST 3000 and LX 2000 LogLogic shipped us two appliances for testing: an ST 3000 and an LX 2000.



The LX is a classic log-aggregation, monitoring and storage product, while the ST provides a platform for large (terabyte-size) unaltered logs that are easily text-searchable. The ST and LX devices both receive and store data, but the ST is designed to manage larger data sets. The two products are manageable through the same UI, however, so interaction was seamless; we typically didn't even notice which unit we were accessing. LogLogic positions the two appliances as a bundle for organizations that need ad hoc querying abilities on larger data sets.

Both LogLogic's strengths and weaknesses share a single trait: simplicity. On the positive side, we had the products up and running in a matter of minutes. Similar to High Tower's product, our LogLogic appliances attempted to identify the log sources it received, which made the configuration process painless. The UI is Web-based, and we found it a bit cluttered at times, but for general reporting and ad hoc querying it's adequate.

Configuration and user setup were easy—we even configured remote NAS devices for additional storage capacity in no time. LogLogic also has a basic but incredibly useful ad hoc querying tool. You can enter a text string, for example, and it will search its entire data set—regardless of data origin or device type—for that string. This may sound like a simple enough feature, but we were surprised at how few products can accomplish this easily. LogLogic's capabilities had us using the product more than the others for tasks such as username lookups.

On the negative side, LogLogic seems more focused on simply managing and storing logs, and less with real-time monitoring and analysis. The product does have a basic dashboard for

monitoring event-per-second loads, system health and general statistics on log transportation, but it lacks the real-time correlation rule sets found in ArcSight ESM and QRadar, making it less security-analyst- and incident-response-friendly. Bottom line, you probably aren't going to use it to monitor the health of your security posture; it's more suited to monitoring the health of your logging infrastructure. Now, this log emphasis isn't necessarily a bad thing—log transportation, storage and querying capabilities are all necessary components of security information management. However, organizations need to be mindful of their requirements. For some, LogLogic's products might meet their needs. But those that want a more SOC-minded toolset with advanced correlation rules may need to invest in a third-party analysis tool in addition to LogLogic to meet their goals.

**SenSage Enterprise Security Analytics** SenSage, formally known as Addamark, got its start in the SIM arena by providing a powerful back-end replacement for conventional RDBMS technology. Originally, SenSage partnered with SIM vendors to bundle its technology, but over the past 18 months the company has begun integrating features into its ESA product that are squarely in the area of SIM. SenSage shipped us a whopping six Intel multiprocessor sys-

tems running Linux on which to evaluate ESA. It offers a software-only delivery model, but the company clearly wanted us to evaluate ESA on a specific set of hardware.

SenSage's claim to fame is its proprietary database-like architecture that does away with much of the unnecessary overhead found in most SIM apps. The rough idea is that, by structuring and indexing data a specific way, you can scale your online SIM database to multiple terabytes without being plagued by conventional RDBMS performance ceilings. In speaking to organizations that have crossed the multi-terabyte threshold in their SIM deployments, we're convinced that performance can become a problem with really large data sets. However, we weren't able to test this.

Although we would love to put the data-volume issue to the test in future review, we humbly submit that data management is but one facet of SIM; if the reporting, UI and real-time analysis capabilities aren't there, query times are going to be a secondary concern. Unfortunately, these other areas are precisely where ESA needs work.

SenSage has most of the important pieces of a SIM product in various stages of development, but it reminded us of just how inadequate offerings from SIM vendors were two to three years ago. For example, the installation process was far from painless and took several days to complete, we had a few glitches with

## SIM Product Features

	ArcSight ESM	High Tower Software Security Event Manager	LogLogic ST 3000 and LX 2000	Network Intelligence enVision	OpenService Security Threat Manager 3.5	Q1 Labs QRadar	SenSage Enterprise Security Analytics	Symantec Security Information Manager 9500
<b>Asset weighting functionality</b>	Y	Y	N	Y	Y	Y	Y	Y
<b>Built-in real-time correlation rules</b>	Y	Y	N	Y	Y	Y	Y	Y
<b>User-definable real-time correlation rules</b>	Y	N	N	Y	Y	Y	Y	Y
<b>Input filters for event viewers</b>	Y	N	Y	N	Y	Y	N	Y
<b>Real-time event viewer for correlated/noncorrelated events</b>	Y/Y	Y/N	N/Y	Y/Y	Y/N	Y/N	Y/N	Y/Y
<b>External IP threat feeds</b>	N	N	N	N	N	N	N	Y
<b>Built-in database-management tools</b>	Y	N	N	N	Y	Y	N	N
<b>SIM time-stamp correction</b>	Y	Y	Y	Y	N	Y	N	N
<b>Product delivery method</b>	Software	Appliance	Appliance	Appliance	Software	Both	Software	Appliance
<b>Autotable rotation/purging</b>	Y	Y	Y	Y	Y	Y	N	Y
<b>Vulnerability scanner support</b>	Y	N	N	Y	Y	N	N	N
<b>Bundled database</b>	None	Oracle	MySQL	Proprietary	MySQL	MySQL	Proprietary	N/A
<b>Supported back-end databases</b>	Oracle	N/A	N/A	N/A	N/A	N/A	N/A	N/A
<b>Data transport</b>								
Batch queuing	Y	N	N	N	N	N	N	N
Bandwidth throttling	Y	N	N	N	N	N	N	N
Encrypted transport	Y	Y	Y	N	Y	Y	N	N
<b>Pricing</b>								
Starts at	\$25,000	\$35,000	\$24,999	\$22,000	\$43,000	\$56,000	\$70,000	\$59,000
As tested	\$67,500	\$60,000	\$79,999	\$64,900	\$78,000	\$65,366	\$200,000	\$74,000
Additional hardware estimates	\$15,000	None	None	None	\$9,000	None	\$30,000	None

Y=Yes, N=No

services dying during our testing and requiring restarts, and the UI is closer to the Microsoft Office toolbar than an integrated SIM suite. Bringing new devices and parsers online was painful, too, requiring us to copy and edit files on the file system to get everything working properly. None of these are show-stopping faults, but we found the product seriously rough around the edges. It's also stunningly expensive, with an as-tested price about three times greater than what rivals charge. It is one of the least-mature SIM products in the space plus the highest entry-level price—not a winning combination.

To the company's credit, our support team solved every problem we tossed at it, including a rapid turnaround on an arpwatch parser and alerting mechanism we wanted to put in place.

Perhaps for some shops SenSage's scalable back end will balance its price and lack of polish, but we suspect the company is going to need to beef up the rest of the product's features to remain competitive.

### OpenService Security Threat Manager 3.5 OpenService

**C** primarily uses a software delivery model, though it told us it has provided turnkey packages in the past.

We tested STM 3.5 using stock Windows 2003 builds on our Dell 1850 servers. STM has a few deployment options, but for our review we decided to put the event collectors and correlation engine on the same platform. Initial installation was a little confusing, but it didn't take us more than a few hours to get the base system up and running. Bringing our device data feeds online was more complicated, however. In fact, we didn't find very many operations in STM intuitive.

OpenService has a lot of the basics covered: STM supports a wide range of devices; it has reporting tools and a correlation rule capability; and it provides built-in database management tools to import, export and prune stored data—a feature a number of the other products lack. It also possesses one of the more advanced role-based authorization systems

we've seen, letting organizations tier access to the SIM infrastructure. Access can be controlled for editing configuration files, restricting which device classes can be viewed, viewing alerts, managing the database and more.

What STM doesn't have, however, is a friendly and intuitive UI and an easy way to manage everything. The UI is HTML- and Java-based and uses a tree-based navigation method. Similar to QRadar's interface, STM's presents a lot of actual event data using HTML-based tables. Unfortunately, we found the layout prohibitively confusing. STM also lacks a native syslog collection service; OpenService expected us to install a third-party service to get things running, a state of affairs we find shocking in 2006.

STM is part of a larger OpenService network management framework, but testing integration with that framework was beyond the scope of this article. Perhaps customers looking to have more closely linked network and security operations teams will find STM more useful than we did, but we think OpenService would be wise to modernize much of its product. **NWC**

**GREG SHIPLEY** is the CTO of Neohapsis, an information security consultancy and enterprise IT product-testing lab.



ArcSight, a leader in Enterprise Security Management (ESM), provides real-time threat management and compliance reporting yielding actionable insights into your security data. By comprehensively collecting, analyzing and managing security data, ArcSight ESM™ enables enterprises, government organizations and managed security service providers to centrally manage information risk more efficiently. ArcSight's customer base includes leading worldwide companies across many verticals—and more than 20 of the largest U.S. federal agencies. For more information, visit [www.arcsight.com](http://www.arcsight.com).

## IS RDBMS BAD IN THE SIM WORLD?

One of the industry's ivory tower debates might see the light of day soon as SIM deployments in large organizations become more common. The discussion goes something like this: Most SIM implementations are a write-once, read-many scenario; logs and event data are written to a storage medium of some sort, then searched and referenced over time.

Conventional relational databases, on the other hand, are built with all sorts of other applications in mind—multiple reads and writes, record locking, and journalized transactions are just some of the features used in transaction scenar-

ios, but not necessarily with SIM.

Here's the debate: One camp claims that a simplified, purpose-built approach is needed because response times go down the tubes once data sets become large. The other camp believes that with good data management and enough hardware, commercial off-the-shelf RDBMSs should do just fine. Based on our observations of current, large-scale production SIM deployments, performance seems to become a problem only after databases get to multiple terabyte territory. For some organizations, that could be just one or two months'

worth of data, making performance a valid operational concern. For others, it won't become an issue until years of data are stored. And even in that scenario, how often do you go back and perform time-critical searches on ancient data sets?

We see merits in both schools of thought, but more testing and time is needed before weighing in conclusively. One thing is for sure: SIM data volumes are only going to increase. This is a debate worth watching.