National Information Assurance Partnership

Common Criteria Evaluation and Validation Scheme Validation Report

Securify

SecurVantage™ Version 5.0

Report Number: CCEVS-VR-06-0012
Dated: 24 February 2006
Version: 1.0

National Institute of Standards and Technology
Information Technology Laboratory
100 Bureau Drive
Gaithersburg, MD 20899

National Security Agency
Information Assurance Directorate
9800 Savage Road STE 6740
Fort George G. Meade, MD 20755-6740

TABLE OF CONTENTS

1. Executive Summary

This report documents the National Information Assurance Partnership (NIAP) validator's assessment of the evaluation of the SecurVantage™ Version 5.0, a product of the Securify Inc., Cupertino, CA 95014. Securify SecurVantage™ is a system that enables customers to define a network security policy (typically describing the permitted network operations), monitor networks for compliance with that policy, and produce relevant network operational information (such as events that fail to comply with the policy). Securify's SecurVantage™, version 3.1 had been evaluated in January 2004 at Evaluation Assurance Level (EAL) 2.

The evaluation of the Securify's SecurVantage™, version 5.0 at EAL 3, was begun in April 2005 and completed in February 2006. The product consists of a tool for network security policy development and security analysis (Studio), a real-time monitoring system to continuously verify conformance to those security policies (Monitor), and an enterprise management console (optional) that can merge multiple monitoring points into a single visual display (Enterprise). The (optional) report generating system, Enterprise Reporting Gateway supports quantitative network and application trend reporting.

The Target of Evaluation (TOE) includes those components developed by Securify, and not third-party components such as hardware and operating systems. The evaluation examined the threat of unauthorized users gaining control of the TOE, of attackers evading the monitoring implemented by the TOE, and of users raising their privileges in an unauthorized way. It is assumed that the TOE hardware and software will be located within controlled access facilities, preventing unauthorized physical access and protecting the TOE from unauthorized physical modification.

The evaluation was performed by the CygnaCom Common Criteria Testing Laboratory (CCTL), and was completed during February 2006. The information in this report is derived from the Evaluation Technical Report (ETR) and associated test reports, all written by the CygnaCom CCTL. The evaluation team determined that the product is Common Criteria version 2.2 [CC] Part 2 and Part 3 conformant, and meets the assurance requirements of EAL 3 from the Common Methodology for Information Technology Security Evaluation, Version 2.2, Part 2: Evaluation Methodology [CEM]. The product is not conformant with any published Protection Profiles, but rather is targeted to satisfying specific security objectives while countering specific threats.

The evaluation and validation were consistent with National Information Assurance Partnership (NIAP) Common Criteria Evaluation and Validation Scheme (CCEVS) best practices as described within CCEVS Publication #3 [CCEVS_PUB 3] and Publication #4 [CCEVS_PUB 4]. The Security Target (ST) for Securify SecurVantage is contained within the document Security Target for Securify SecurVantage Version 5.0.

The information contained in this Validation Report is not an endorsement of Securify's SecurVantage™ by any agency of the U.S. Government and no warranty of the product is either expressed or implied.

1.1 Evaluation Highlights

Dates of Evaluation:	April 2005 – February 2006
Evaluated Product	SecurVantage™, version 5.0
Developer:	Securify Inc., 20425 Stevens Creek Blvd., 2nd floor, Cupertino, CA 95014 (http://www.securify.com)
CCTL:	CygnaCom
Lead Evaluator:	Nithya Rachamadugu
Evaluation Class:	EAL 3
PPs Claimed:	None.
Validation Body:	NIAP Common Criteria Evaluation and Validation Scheme
Version of CC:	Common Criteria version 2.2, January 2004
Version of CEM:	Common Evaluation Methodology 2.2, January 2004

2. Product Identification

The Target of Evaluation (TOE) is Securify SecurVantage™, version 5.0. It consists of the following major components:

- Studio 5.0
- Monitor 5.0 or Monitor LE 5.0
- Enterprise 5.0 (optional)
- Enterprise Reporting Gateway 5.0 (optional)

The Monitor component can be configured in two different ways, depending on the processing speed required. High bandwidth solutions use Monitor 5.0, which has two subcomponents (Monitor SM and Monitor Harvester) each running on a different processor. For low bandwidth solutions, Monitor 5.0 LE is deployed, which uses the same software but uses a single hardware system that is shared between the two subcomponents.

The product must run on top of hardware and an operating system, all of which are part of the environment (and thus not evaluated). Monitor, Enterprise, and Enterprise Reporting Gateway run atop Red Hat Linux 7.2 (patched) and an x86 system. Although it is also sold separately, in the evaluated configuration the hardware and operating system for Monitor and Enterprise are provided by the vendor directly to the customer, along with a CD-ROM permitting local re-installation. Studio is a user interface component and runs on Microsoft Windows XP on an x86

hardware platform. Additional interface is via a web browser, and the system depends on third-party encryption libraries, neither of which are part of the TOE.

3. Security Policy

The TOE, with support from its IT environment, provides the following security functions:

- Auditing,
- Access Control,
- User Identification and Authentication
- Security Management

The primary purpose of the TOE is to permit administrators to define a network security (SecurVantage) policy, then monitor networks and report deviations from that policy. A SecurVantage Policy is a set of rules that describe the expected behavior of the systems within a network. Network objects represent systems. A network object can be one or many IP addresses. Each rule in the Policy describes how the system will log a network transaction between two network objects. All network transactions are logged and represented as an event. Each event represents the information contained in the headers of the actual packets within the network transaction. In SecurVantage, an output of the policy engine is created when network traffic is evaluated against a policy. A network event is a summary of the set of protocol events that make up a complete application level session on the network. For example, viewing a Web page creates a network event that summarizes the underlying Internet Protocol (IP) association, Transmission Control Protocol (TCP) connection and Hypertext Transfer Protocol (HTTP) Get protocol events.

The policy assigns by default a severity to every event, such that all events are logged by default. These default values can be changed by the user of the system to accommodate specific security policies. A severity is one of the following options: Critical, High, Medium, Warning, Monitor, Informational, or Ok. All events other than Ok are fully logged in the system down to the protocol details level (source and target network object name, ip addresses, protocols, src port, dst port, tcp flags, udp association, etc). Events that have a severity value of "ok" are only logged at a summary level (source and target network object name and service name). Events logged as critical are also called alerts and copied to a separate alert table. Alerts can trigger Simple Mail Transport Protocol (SMTP) and Simple Network Management Protocol (SNMP) messages to other management systems.

Effort was particularly expended to ensure that attackers could not easily inhibit or circumvent this monitoring. A given monitor can only implement one network security policy at a time. Multiple monitors may be part of a domain (or security zone as referenced in the administrator

guides), which shares a common network security policy. Enterprise (when used) can support multiple domains (or security zones).

The TOE supports various user roles. Every user account is assigned one or more roles. The privileges granted to a user are the union of the privileges of that user's role(s). All roles (and their privileges) are simultaneously active for a given user. For example, a user with the role "operator" and/or "analyst" (and no other role) cannot upload a new network security policy (i.e., cannot change the policy). The various roles, and their privileges, are shown in the following two tables:

SecurVantage™ User Access Policy (from Security Target)

User Access Policy: Roles/Subjects (Monitor and Enterprise)					
Objects	Operator	Analyst	Developer	SV Manager	Account Manager
Event Data	View	View	View		
Machines	View Status	View Status	View Status Start/Restart Stop	View Status Start/Restart Stop Configure	View Status
DMEs		Download	Download		
User Access					Manage
Policy History	View	View	View		
Policies		Extract	Upload Revert Extract		
Alerts	Manage	Manage	Manage		
Application Logs				View	
User Logs					View

SecurVantage™ ER Gateway User Access Policy (from Security Target)

ER User Access Policy: Roles/Subjects		
Objects	ER SV Manager	ER Account Manager
Machines	View Status Start/Restart Stop Configure	View Status
User Access		Manage
Application Logs	View	

ER User Access Policy: Roles/Subjects		
Objects	ER SV Manager	ER Account Manager
User Logs		View

Details, including definitions of these objects, are given in the Security Target.

4. Assumptions and Clarification of Scope

4.1 Usage Assumptions

The following usage assumptions were made for the TOE.

It is assumed that administrators will have a strong understanding of the TOE, networking technology, and the network(s) they are monitoring. Fundamentally, the TOE allows an administrator to identify a policy of "normal" behavior, and the TOE will then report on all actions not corresponding to the policy. A knowledgeable administrator who creates a well-defined policy may find this TOE to be extremely effective at reporting just the events that need reporting. However, a poorly-defined policy (created by an administrator with insufficient understanding) may cause the TOE to report a voluminous number of unimportant events, and/or cause the TOE to omit events that were important to report. This is not a fault of the TOE implementation; it is fundamental to the nature of its approach. Any TOE is best used by a knowledgeable administrator, but this TOE in particular requires a good administrator for effective use. Administrators should obtain training before use; the vendor makes such training available. Many deployments may want to ensure there are at least two trained administrators, to enable discussions of policy and to ensure continuous service if an administrator becomes unavailable. Administrator training and an understanding of the network being monitored are critical for effective and efficient use of this TOE.

At a more fundamental level, this TOE requires that it be possible to (eventually) determine the expected or permitted activities on the monitored networks, so that this information can be captured as a security policy. If all actions are permitted by all network components, the TOE's ability to compare actions with expected actions is far less valuable. A pre-existing written security policy, while very helpful, is not required; the security policy can be developed over time, starting with a more general policy and then repeatedly refining it. The TOE can also be used as a monitoring tool, so that actual network activity can guide formulation of the security policy. For nearly all real-life circumstances this is not a restriction. Most of today's networks *do* have a set of expected activities that is a small subset of all possible activities.

The TOE is normally used by plugging it into a switch's (Spanning Tree Protocol) SPAN port. Typically SPAN ports only report the packets that *cross* the switch, and not network packets that appear on a network but do not cross it. This is an aspect of the environment, not the TOE itself; the TOE can only log what's reported to it. Since this is not an issue of the TOE itself, and many customers would expect this behavior anyway, it is not considered a vulnerability. However, administrators will need to configure their network and/or where they connect Monitor(s) so that what they wish to monitor can actually be monitored.

Users of the system are trusted with the privileges they have been granted, and it is presumed that authorized users will not misuse their privileges. For example, operators are trusted with the ability to view event data (which would give operators insight into all network activity) and analysts are granted the ability to download DME (a Securify Proprietary Format) files and extract policies (which would give analysts the ability to download summaries of network activities and know exactly what policy is being checked). However, as clearly noted above, users of the system are not fully trusted with all privileges. The TOE specifically works to prevent authorized users from gaining additional unauthorized privileges.

4.2 Environmental Assumptions

As stated in the ST, the hardware, operating system, and third-party cryptographic libraries are not included in the evaluation. As part of installation, the operating system used by Monitor and Enterprise (Red Hat Linux) is installed in a way that limits its functionality. While this should help, for purposes of evaluation it is assumed that this operating system is secure in its environment. Note in particular that the operating system running Studio (Microsoft Windows) must be secured separately, and that no attempt is made by the product to ensure this. The TOE uses encryption to protect data between its major components, but since this encryption is performed in the environment (not in the TOE) it is not further considered here.

As clearly stated in the ST, the TOE only records IPv4 with normal Ethernet framing. Other kinds of data, particularly IPv6 and Ethernet jumbo frames (jumbograms), are not recorded. If all relevant network traffic is to be monitored, systems must be configured to reject these unrecorded packets. Many systems do this by default, so this is a plausible environmental restriction. If there are concerns that cooperating end-systems may surreptitiously send data between each other using other kinds of packets, then the network infrastructure should be configured to actively inhibit this kind of traffic. A particularly effective approach to doing this would be to insert network packet scrubbers that enforced these limits and regularized packets for monitoring purposes.

The TOE is capable of using the Domain Name Service (DNS) for translating IP addresses back to machine names. By default, this capability is disabled. Enabling this capability can aid administrators by giving them simple names instead of IP addresses. However, these name values are dependent on the security of DNS itself. Subversion of the DNS service could

provide incorrect names. Also, attackers may control DNS services of other domains (legitimately or not). Thus, the names provided by DNS could be misleading. Note that the TOE does not use DNS for security decisions—this data is purely informational. The TOE's Studio component does include an ability to import DNS zone data from a file; ensuring that this zone data file is correct is outside the scope of the TOE. Administrators are warned about these issues in the installation guidance.

The TOE is capable of using the Network Time Protocol (NTP) for keeping time values correct. By default, this capability is disabled. NTP can be convenient for accurately keeping time values current. However, it is difficult to secure. An attacker that sends malicious NTP reports, or takes over a relevant NTP server, could cause the timestamps of events to be incorrect (impacting any Monitor report). This could also negatively impact attempts to merge data from multiple Monitors (as Enterprise does). Administrators are warned about this in the installation guidance (in the Administrator Addendum).

The TOE hardware and software must be located within controlled access facilities, preventing unauthorized physical access and protecting the TOE from unauthorized physical modification.

4.3 Clarification of Scope

All evaluations (and all products) have limitations, as well as potential misconceptions that need clarifying. This text covers some of the more important limitations and clarifications of this evaluation. Note that:

1. As with any evaluation, this evaluation only shows that the evaluated configuration meets the security claims made, with a certain level of assurance (EAL 3 in this case).
2. This evaluation covers SecurVantage Version 5.0, not the later version 5.X series. Thus, these evaluation results do not automatically apply to version 5.X, and in particular additions in version 5.X (such as the Nessus security scanner) have not been considered by this evaluation.
3. As with all EAL 3 evaluations, this evaluation did not specifically search for, nor seriously attempt to counter, vulnerabilities that were not "obvious" (as this term is defined in the CC and CEM) or "vulnerabilities" to objectives not claimed in the ST. In particular, the evaluation does not claim a resistance of the TOE to denial-of-service attacks, where an attacker intentionally causes a large load of traffic to mask different nefarious activities. This is a fundamental limitation of any passive monitoring tool. However, the TOE *does* log when its logging rates have been exceeded. Thus, although an attacker might be able to hide a nefarious action by overwhelming the TOE, the TOE *will* log when there was an opportunity to do so. Customers concerned about attackers who would use resource overwhelming attacks to hide other actions may consider purchasing the high-bandwidth Monitor, employing multiple Monitors, and/or employing network rate limiting components to slow network traffic to a monitorable rate.

4. Ideally, the monitoring system should log *exactly* what the receiving application would see. However, for a variety of reasons this is impossible for a passive monitoring system to do perfectly. In particular, when the TOE reassembles IP fragments, the TOE may reassemble them differently than the receiving system and thus report different results. This is a fundamental problem with all network monitoring systems, particularly passive monitoring systems ones like this. Since this case is a fundamental limitation of the technology, and would be expected by knowledgeable administrators anyway, this was considered acceptable. Customers who find this undesirable should counter this problem by inserting "network packet scrubbers" in their environment in such a way that the monitor and end-user receiving systems will see exactly the same data.
5. Network events can be a set of IP packets, not just one (e.g., initiating an HTTP request). If an event is halfway, but never actually completed, the TOE times out. However, the result of this time-out is logged, and the time-outs are longer than the standards required of the receiving systems. This again raises the concern that the receiving systems and the TOE may have a slightly different view of what is happening, but since there is some logging no matter what the attacker does this was determined to be acceptable.
6. The purpose of the TOE is to report violations of policy, not identify network covert channels between cooperating parties. For example, two systems A and B could surreptitiously communicate with each other by communicating with a system C in accordance with a security policy, and then by monitoring each others' communication with C, extract the separate data intended for each other. Alternatively, systems A and B could arrange for data or processing at C to be an indirect channel between A and B. Steganographic tricks combined with error-correcting codes could make such hidden data particularly hard to identify. Detecting such channels is *not* the purpose of this TOE. The purpose of this TOE is to watch ordinary protocols and report direct violations according to a customer-defined security policy. Customers worried about these covert channel problems should consider redesigning their networks (e.g., to completely isolate the systems) or other measures.
7. The TOE does not log every byte of every packet involved. In many of its intended environments, this would be an extremely stressing requirement, would severely limit the length of time the logged information could be stored, and is unnecessary. Instead, the TOE stores a summary of every network event, with more information on events that are not ranked "ok." This information is sufficient for its intended purpose.
8. Network event data is stored at the Monitor that recorded it. This includes the criticality of the event, which is treated as a constant (it is assigned using the policy active at the time the event was recorded). Network event data can be retrieved as a DME file (if the user is permitted to do so); Studio can then be run locally to recompute network event severities using a different policy. These recomputed severities do not change the severity recorded by the relevant Monitor.
9. SecurVantage™ Monitor audits IPv4 packets when transported over Ethernet frames with length less than 1518 bytes. Notice that this includes neither IPv6 packet nor Ethernet

Jumbo frames. Monitor does not support these types of packets. Such frames are discarded without logging and policy evaluation.

The product is not conformant with any published Protection Profiles, but rather is targeted to satisfying specific organizational security policies while countering specific threats. The ST provides additional information on the assumptions made and the threats countered.

5. Architectural Information

SecurVantage™ consists of four major components:

1. SecurVantage™ Studio: provides a management interface that allows for the authoring of network security policy at multiple levels.
2. SecurVantage™ Monitor: captures and evaluates monitored network traffic according to the security policy
3. SecurVantage™ Enterprise: combines the information from multiple monitoring points into a single, real-time monitoring and management console.
4. SecurVantage™ Enterprise Reporting Gateway: provides quantitative network and application trend reporting.

Users use Studio to define the network security policy using a proprietary policy language. This security policy defines the "correct" behavior of the network(s) being monitored.

Users may use a web browser to communicate with Monitor or Enterprise over an encrypted, Secure Sockets Layer (SSL), link. When initially connecting to Monitor or Enterprise, users are presented with a new self-signed server certificate, which they can verify by comparing these certificates to the values generated during Monitor/Enterprise installation. Users are first authenticated (using username and password sent over the encrypted link), and depending on their roles users may change data (such as uploading a new policy defined using Studio) as well as receive data (such as event data). Note that the web browser and third-party encryption libraries are outside the TOE.

Studio also supports query-only direct access to Monitor or Enterprise. When using Studio in this way, Studio also uses an encrypted SSL link (to the same SSL/HTTP port). This access approach supports the username/password pair (it is the same mechanism as above), but in addition it also permits users to set up client-side certificates. Note that this direct access only permits query operations, so only users with the role of operator, analyst, or developer, can usefully use this access approach. This access approach cannot be used to modify information; in particular it cannot be used to modify policies.

SecurVantage™ Monitor captures and evaluates in real time the packets flowing through the network at all levels of the protocol stack. It then makes decisions on whether the traffic is consistent with the policy specification. The result is a set of "network events" with each event including an attribute termed "criticality" in the ST.

If an Enterprise system is deployed, Enterprise copies information from the Monitors connected to Enterprise and aggregates them into a local database. This database is accessible through the web interface for a period of 48 hours. The Enterprise serves also as a conduit to the Monitors' databases when detailed information is requested by Studio application. Enterprise can also deploy policies to multiple Monitors.

When Enterprise is deployed, Enterprise and the Monitor(s) communicate using SSL. For this communication, certificates are initially verified the first time the Enterprise and Monitor components communicate.

An Enterprise Reporting (ER) system is composed of an ER Gateway and an ER Warehouse—each installed on separate machines. The ER Warehouse is not part of this evaluation. The ER Gateway collects data from one or more Enterprise systems. Every hour, the ER Gateway prepares the data it has collected, and then inserts the data into the ER Warehouse. Data can remain in the ER Gateway for up to 36 hours, depending on the volume of data. Administrators can communicate with ER Gateway via web interface using SSL.

Figure 1 shows a typical deployment of SecurVantage™, although SecurVantage™ Monitor can be placed anywhere on the network. It does not necessarily have to be on its own sub-network and does not have to be connected through a switch. Typically SecurVantage™ Monitor is connected to the SPAN port of a switch where there is traffic relevant to the policy. However, there are no assumptions about the source of the traffic.

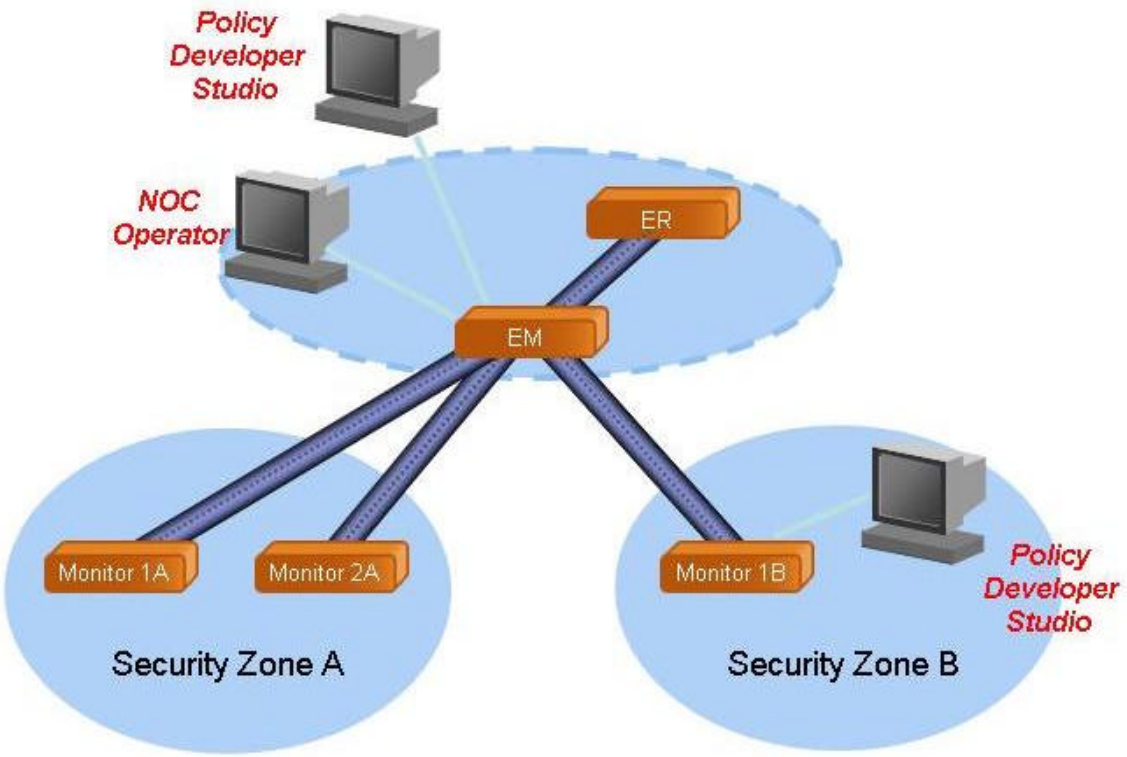

Figure 1: Typical SecurVantage™ Deployment

SecurVantage™ consists of the policy development and analysis environment coupled with the monitoring system and the enterprise management system. Figure 2 shows the System Architecture.

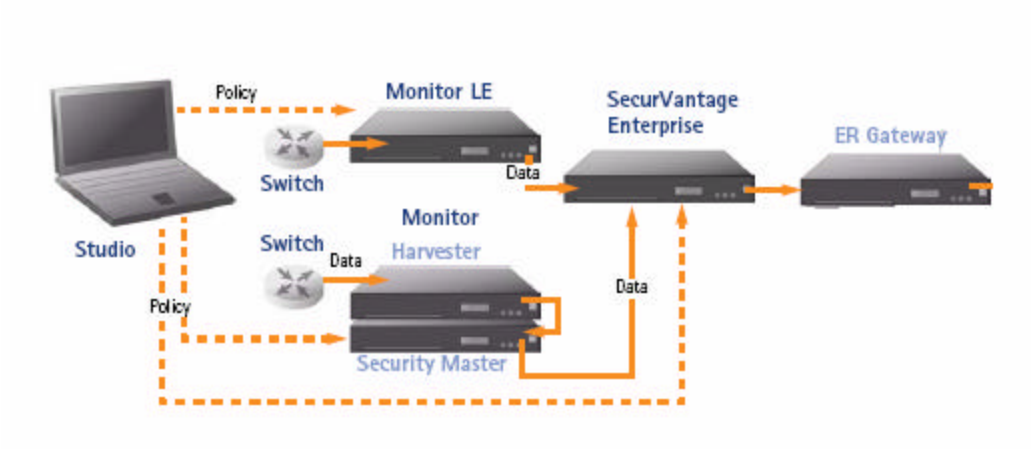

Figure 2: SecurVantageTM System Architecture

See the Security Target section 2.2 for additional discussion of the TOE's architecture, including more information on the tasks performed by each major component.

6. Documentation

The documentation provided with the product to customers is as follows:

- Securify SecurVantage 5.0 Installation Guide. January 2005.
- Securify SecurVantage 5.0 Release Notes. January 2005.
- Securify SecurVantage 5.0 Common Criteria Addendum. January 2006
- Securify SecurVantage 5.0 Enterprise Reporting Operations Guide. December 2004
- Securify SecurVantage 5.0 Web Interface. December 2004
- Securify SecurVantage 5.0 Studio User Guide. December 2004
- Securify SecurVantage 5.0 Administrator Addendum. January 2006

7. IT Product Testing

7.1 Examination of Vendor Tests

The vendor provided test plans, procedures, test results and a test coverage document. The security testing was driven by the SecurVantage "Test Matrix," a spreadsheet of tests grouped into three sections: (1) IT Security functions related (as section 6 in the ST), (2) tests for Delivery, Installation and Configuration, and (3) other tests. The IT Security functions related tests are indexed by the functions as identified in the section 6 of the ST and further grouped into: Studio, Monitor, Enterprise and Reporting Gateway.

Each test was performed by one of the following "test tools":

1. Test Matrix: This is a document (supporting the spreadsheet) that specifies step by step how to perform a specific test. The document is indexed by a test case number.
2. Sentinel: This is an automated tool for testing web applications. The test case is usually mapped to a signature of the sentinel test.
3. By Inspection: Guidance necessary to perform the test.
4. By code inspection: Tests that can only be verified through a source code inspection

For purposes of testing, a special test harness was used. Normally, the system is connected into a switch's SPAN port. However, testing the system by directly connecting to a SPAN port and then generating data from multiple different networks would result in repeated tests not generating identical inputs to the TOE (due to different interleaving). Thus, for testing purposes, instead of connecting to a switch's SPAN port, the TOE was directly connected to systems which replay data previously captured from a SPAN port. The evaluator and validator examined this configuration and were satisfied that the test harness produced the same inputs (as seen by the TOE) as the original network whose traffic had been captured (including Ethernet MAC addresses).

The evaluator determined that the vendor tested (at a high level) most security-relevant aspects of the product. The evaluator determined that the developer's tests were sound in their approach. The test document provided the configuration of the test hardware and software, the objective for each of the tests, and test procedures. The information provided was adequate to be able to reproduce the tests. The evaluators determined that the developer's approach to testing the TSF was appropriate for this EAL 3 evaluation.

7.2 Evaluator Independent Tests

The evaluator performed the tests at the developer's site using the equipment provided by the developer. The tests were performed in a configuration representing installations of the TOE which contained two Enterprise Managers; one managing one Monitor and another managing two monitors containing both (Monitor and LE) versions of the Monitor, one Studio and one Enterprise Gateway. The Enterprise Managers, Enterprise Gateway and the Studio were connected via separate network to the Monitors. Though this is not an imposed restriction, this is the most common field installation configuration.

The evaluator installed the TOE using the installation procedures. About 30% of the developer tests were repeated. The evaluator used the developer's automated tool Sentinel for some of the tests. The evaluator repeated some tests manually to gain confidence in the tool. About 25% of the developer's manual tests were also repeated. The developer provided a traffic generator tool that replayed the traffic from a previous session. This provided a varied sample of the traffic and helped to regulate the speed of the traffic as desired.

Functions that were deemed critical to the operation of the TOE components (for example. DME creation, policy pushing, creating and revoking of users, critical event viewing), and all new operational scenarios from ER Gateway were chosen as the basis for the evaluator's tests.

Critical messages were generated and the SNMP, SMTP message generation was verified for critical alerts. For the test repeated from the developer suites, the evaluator examined the test results and found them to be matching those of the developer. Any mismatches were purely due to data related inconsistencies. The overall verdict of the evaluator testing is that the TOE components perform the security functions in accordance with those specified in the ST and the developer's test results match those of the evaluators.

7.3 Strength of Function

Even if a TOE security function cannot be bypassed, deactivated, or corrupted, it may still be possible to defeat it because there is a vulnerability in the concept of its underlying security mechanisms. For those functions a qualification of their security behavior can be made using the results of a quantitative or statistical analysis of the security behavior of these mechanisms and the effort required to overcome them. The qualification is made in the form of the strength of TOE security function claim.

The security of this TOE depends on the strength of the passwords used to access Monitor, Enterprise, and Enterprise Reporting Gateway. The administrator guidance includes the following password policy, which is enforced at the TOE's web interface:

- Minimum of 8 characters in the password (maximum 64 characters).
- At least one lower case character.
- At least one upper case character
- At least one numeric character.

A strength-of-function analysis took these password requirements (as well as other information) and justified a ranking of SOF-basic, which effectively requires resistance to password guessing attacks of at least one day. The overall SOF requirement for the TOE made in ST is expressed as an SOF rating, SOF-basic. Thus, the TOE meets the ST requirements.

7.4 Vulnerability Analysis

Vulnerability analysis is a process for identifying potential vulnerabilities and determining whether potential vulnerabilities identified throughout the evaluation process could allow users to violate the TSP. See the CC and CEM for additional information on the requirements for an EAL 3 evaluation.

The vendor searched for publicly known vulnerabilities specifically related to SecurVantage, as well as publicly known vulnerabilities in the third-party products used by SecurVantage (and are in the environment of the evaluation). This search included searches in the information of Security Focus (www.securityfocus.com), Bugtraq (through Security Focus), Packetstorm (www.packetstormsecurity.org) and CVE (cve.mitre.org). No publicly-known vulnerabilities specific to SecurVantage were found. The vendor also considered obvious vulnerabilities associated with the Enterprise and Monitor web interfaces and process of collecting, storing and presenting network traffic.

Known vulnerabilities in the IT environment could also be exploited to bypass the TOE's security policies, but these are normally outside the scope of the evaluation. In particular, the customer is expected to install the latest security critical patches to the Windows XP operating system. Under unusual circumstances a patch to TOE may also be required to address compatibility issues with a specific operating system patch. The customer is advised to check the vendor's support web site for any restrictions on specific patches to components of the IT environment. The evaluation testing activities were performed using Windows XP Professional SP2 with the latest security-critical patches installed.

The TOE is bundled with several significant IT environment components that are maintained by the vendor. Publicly available patches to or newer versions of these components cannot be installed independently by end-users. These components, while not part of the TOE, were included in the search of obvious vulnerabilities that was performed as part of the AVA_VLA.1 work units. In particular, the known vulnerabilities in the Red Hat Linux 7.2 operating system's RPMs were examined, and were either countered or shown to be unlikely to be exploitable in the TOE's configuration.

The assumed level of expertise of the attacker is unsophisticated, with access to only standard equipment and public information about the product. The specific threats that the TOE is designed to counter are listed in section 3.2 of the ST. The evaluator determined that the product met the criteria of EAL 3 for vulnerability analysis.

8. Evaluated Configuration

The evaluated configuration was configured per the documents listed in section 6 of this report. The system was tested with DNS and NTP services disabled, as is its default (see section 4.2 for a discussion on these services). For additional information on how to securely deploy this TOE (beyond the referenced documentation), see sections 4 and 10 of this report. The following components were evaluated:

Securify SecurVantage™ Version 5.0
Securify SecurVantage™ Studio: 5.0 (V50_CC_7)

Securify SecurVantage™ Monitor (SM): 5.0 (V50_CC_7) with patch "Patch2" (V50_324)
Securify SecurVantage™ Monitor (Harvester): 5.0 (V50_CC_7) with patch "Patch2" (V50_324)
Securify SecurVantage™ Monitor (LE): 5.0 (V50_CC_7) with patch "Patch2" (V50_324)
Securify SecurVantage™ Enterprise: 5.0 (V50_CC_7) with patch "Patch2" (V50_324)
Securify SecurVantage™ Enterprise Reporting Gateway: 5.0 (V50_CC_7) with patch "Patch2" (V50_324)

9. Results of the Evaluation

A verdict for an assurance component is determined by the resulting verdicts assigned to the corresponding evaluator action elements. The evaluation was conducted based upon version 2.2 of the CC and the CEM.

The Evaluation Team assigned a Pass, Fail, or Inconclusive verdict to each work unit of each EAL 3 assurance component. For Fail or Inconclusive work unit verdicts, the Evaluation Team advised the developer of issues requiring resolution or clarification within the evaluation evidence. In this way, the Evaluation Team assigned an overall Pass verdict to the assurance component only when all of the work units for that component had been assigned a Pass verdict.

The Validation Team agreed with the conclusion of the CygnaCom CCTL Evaluation Team, and recommended to CCEVS Management that an EAL 3 certificate rating be issued for the Securify SecurVantage Version 5.0.

The details of the evaluation are recorded in the Evaluation Technical Report (ETR), which is controlled by CygnaCom CCTL. The security assurance requirements are displayed in the following table.

EAL 3 Assurance Requirements

Assurance Class	Assurance Family
ST Evaluation	ASE_DES.1
	ASE_ENV.1
	ASE_INT.1
	ASE_OBJ.1
	ASE_PPC.1
	ASE_REQ.1
	ASE_SRE.1
	ASE_TSS.1
Configuration Management	ACM_CAP.3

Assurance Class	Assurance Family
	ACM_SCP.1
Delivery and Operation	ADO_DEL.1
	ADO_IGS.1
Development	ADV_FSP.1
	ADV_HLD.2
	ADV_RCR.1
Guidance Documents	AGD_ADM.1
	AGD_USR.1
Life Cycle Support	ALC_DVS.1
Tests	ATE_COV.2
	ATE_DPT.1
	ATE_FUN.1
	ATE_IND.2
Vulnerability Assessment	AVA_MSU.1
	AVA_SOF.1
	AVA_VLA.1

10. Validator Comments

As with any evaluation, this evaluation shows that the evaluated configuration meets the security claims made, with a certain level of assurance.

Be sure to note the assumptions and clarifications of scope in section 4 of this report. In particular:

1. Note that this tool is designed to report violations from a customer-defined network security policy. Thus, it is strongly advised that administrators be trained so that they can devise a good network security policy for their environment.
2. Users are added, deleted or edited into Securify servers along with their roles, type of authentication (password or certificate), and the password or hash of the certificate. Users supply the user id and password or the certificate, private key file and private key pass phrase when initiating a connection with a server. Securify enforces server side authentication when user id and passwords are used; but enforces mutual authentication when certificates are used in establishing SSL connection. Certificates are managed by the Managecerts utility, which is described in the Securify SecureVantage Installation Guide Appendix C.

For many purposes, the TOE is useful as it is. However:

1. Consider connecting Enterprise (if present), Studio, and all Monitor(s) using a separate private ("command and control") network solely allocated for this purpose. The evaluated configuration presumed that the network connecting the TOE components was accessible by an attacker, and examined the TOE resistance to attack in that circumstance. No EAL 3 level vulnerabilities were found. However, using a private network can reduce even further the opportunities for an attacker to exploit the TOE or to assail it using denial-of-service attacks. A physically separate network would be even better than a logically separated one.
2. Firewalls may be useful to prohibit certain actions that should be simply prohibited. This would reduce the processing load on the TOE so it can use its processing power to monitor the subtler activity the TOE is capable of monitoring.
3. It may be prudent to pay special attention to the workstation(s) used to run Studio: lock down its operating system to be secure, rigorously maintain its operating system for patches, use it only for SecurVantage-related work (and isolate it administratively for just that purpose), and power down or unplug the workstation when it is not in use. TOE users with lower privileges should not have special privileges to the workstations' operating system used by TOE users with higher privileges, since those special privileges could be exploited to gain control over the other TOE user.
4. Some high-risk environments that decide to use this EAL 3 evaluated TOE may also find it useful to augment this TOE with other tools, since no one tool has all strengths and some tools are especially good complements for this TOE. One useful type of complementary tool would be a network scrubber (with rate limiting), to ensure that the end-systems and the Monitor see exactly the same data and that the data rate is loggable. Another such tool would be specialized intrusion detection systems (IDSs). IDSs could detect attack attempts (even if the communication is allowed by the TOE security policy), while the TOE can detect misuse of the network (including certain kinds of intruder activity) in a way that many IDSs would not detect.
5. Be careful when updating the security policy, since this interferes with capturing network events when the policy is being updated. Avoid updating the security policy at predictable times, and if collecting all network events is critical, consider temporarily disabling the network for a period of time while updating the security policy. As long as security polices are not updated at times known by an attacker, this is more of a theoretical problem than a real one.
6. The safest course is to leave NTP and DNS access disabled. DNS can be enabled, but it's important for users to understand that in some cases an attacker may provide the displayed DNS data. See section 4.2 for a discussion of these issues.

The TOE can detect and report violations of a security policy, but what happens after detection is a decision humans must make. If the goal is to detect authorized users performing certain

prohibited actions, it will be important to ensure that users know (in a general way) what kinds of actions are allowed and what is prohibited. It may be necessary to develop a more detailed human-readable network policy so that authorized users will know what actions they may and may not take.

11. Annexes

None.

12. Security Target

The Security Target is provided separately. It is Version 2.0, dated February 10, 2006.

13. Glossary

The following acronyms are provided for reference:

CC	Common Criteria
CCEL	Common Criteria Evaluation Laboratory
CCEVS	Common Criteria Evaluation and Validation Scheme
CCTL	Common Evaluation Testing Laboratory
CEM	Common Evaluation Methodology
CI	Configuration Items
CSC	Computer Sciences Corporation
DNS	Domain Name Service
DSA	Developer Security Analyst
EAL	Evaluation Assurance Level
EDR	Evaluation Discovery Report
ER	Enterprise Reporting
ETR	Evaluation Technical Report
HTTP	Hypertext Transfer Protocol
IP	Internet Protocol
MRA	Mutual Recognition Arrangement
NIAP	National Information Assurance Program
NIST	National Institute of Science & Technology
NSA	National Security Agency
NTP	Network Time Protocol
OR	Observation Report
PP	Protection Profile
SAR	Security Assurance Requirement

SFR	Security Functional Requirements
SNMP	Simple Network Management Protocol
SMTP	Simple Mail Transport Protocol
SOF	Strength of Function
SPAN	Spanning Tree Protocol
SSL	Secure Sockets layer
ST	Security Target
TCP	Transmission Control Protocol
TCSEC	Trusted Computer Systems Evaluation Criteria
TOE	Target of Evaluation
TSC	TOE Scope of Control
TSF	TOE Security Functions
TSFI	TSF Interface

The following CC terms are provided for reference:

User	Any entity (human user or external IT entity) outside the TOE that interacts with the TOE.
Human user	Any person who interacts with the TOE.
Authorized User	A user that, in accordance with the TOE Security Policy (TSP) may perform an action. (As identified by group membership.)
External IT entity	Any IT product or system, untrusted or trusted, outside of the TOE that interacts with the TOE.
Role	A predefined set of rules establishing the allowed interactions between a user and the TOE.
Identity	A representation (e.g., a string) uniquely identifying an authorized user, which can be either the full or abbreviated name of that user or a pseudonym.
Authentication data	Information used to verify the claimed identity of a user.

14. Bibliography

The evaluation and validation methodology was drawn from the following:

[CC] Common Criteria for Information Technology Security Evaluation, dated January 2004, version 2.2.

[CEM] Common Methodology for Information Technology Security Evaluation - Evaluation Methodology, dated January 2004, version 2.2.

[CCEVS_PUB 3] Common Criteria, Evaluation and Validation Scheme for Information Technology Security, Guidance to Validators of IT Security Evaluations, Scheme Publication #3, Version 1.0, February 2002.

[CCEVS_PUB 4] Common Criteria, Evaluation and Validation Scheme for Information Technology Security, Guidance to Common Criteria Testing Laboratories, Scheme Publication #4, Draft Version 1, March 2001.