SNMP Security Analyzer

The SNMP Security Analyzer is a software tool that assists in the deployment and administration of SNMPv3. If an SNMPv3 manager is unable to communicate with an SNMPv3 agent and it is not clear why, this tool can help find the cause. It also can identify SNMP security exposures that may not be obvious to the network administrators.

Once the SNMP Security Analyzer finds an agent on the network, it conducts extensive testing on that SNMP agent looking for misconfiguations that may prevent SNMPv3 communication, misconfigurations that weaken SNMP security, and known SNMPv3 security vulnerabilities. Once the network-wide analysis is complete, the Security Analyzer automatically generates reports that identify and prioritize the discovered issues.

The SNMP Security Analyzer adds value to the network management toolbox by intensively testing individual SNMP agents, finding and testing all SNMP agents on the network, testing for security issues typically known only by SNMPv3 security experts, then generating network wide reports, all automatically. The level of analysis done by the SNMPv3 Security Analyzer would be difficult, time-consuming, and costly to do manually. The SNMP Security Analyzer looks for problems such as:

- Agents subject to the authentication bypass vulnerability (CERT Vulnerability Note VU#878044). This is a significant security exposure because it allows a non-authenticated user to access the agent. The SNMP Security Analyzer finds devices that have not been updated to remove this vulnerability.
- Multiple SNMPv3 agents with the same SNMP Engine Identifier (snmpEngineID). For security reasons, the snmpEngineID must be unique network-wide in order to generate unique local keys. Duplicate snmpEngineIDs are a common cause of an SNMPv3 manager not being able to talk to an SNMPv3 agent.
- SNMPv3 agents with clock rollbacks. For security, an SNMPv3 manager should reject replies from these agents because it cannot trust the authenticity of the reply. The reply could be a replay attack.
- SNMPv3 agents with clock latches. For security, an SNMPv3 manager should reject replies from these agents. Agents with clock latches rarely happen in practice. The agent must be manually reset to recover.
- SNMP agents that allow non-secure access. This includes SNMPv1 and SNMPv2c agents, which are inherently insecure. It also includes SNMPv3 agents configured for access without authentication or privacy.
- SNMP agents that have incorrect or unused security information configured.

What Security Analyzer Detects	What This Means	Why It Matters	How To Fix
The SNMP Security Analyzer detects whether an SNMPv3 agent is affected by CERT Vulnerability Note VU#878044.	Older SNMPv3 agents may have a security flaw that allows using authentication only, and either a zero-length authentication key or an authentication key containing a specific 1-byte value.	If an older SNMPv3 agent is affected, and an interloper knows an authentication-only user name on that agent, the interloper can access that agent.	Upgrade the vulnerable SNMPv3 agent to a newer version that does not have the vulnerability. A workaround is to enable and use SNMPv3 privacy, which uses a privacy key for encryption in addition to an authentication key.
The SNMP Security Analyzer detects whether two SNMPv3 agents have the same SNMP engine identifier (snmpEngineID).	The SNMPv3 specification requires that each SNMPv3 agent on your network have a unique identifier known as an SNMP engine identifier (snmpEngineID). SNMP Security Analyzer detects whether two SNMPv3 agents on your network violate this requirement.	The snmpEngineID is an input into the process of converting SNMPv3 authentication and privacy passwords provided by a user into security keys. Unique snmpEngineIDs on different SNMPv3 agents insure different security keys on different SNMPv3 agents. If the same snmpEngineID is re-used on SNMPv3 agents, the security keys may be the same on those SNMPv3 agents. Consequently, if one SNMPv3 agent's security keys are compromised, then the security keys on all SNMPv3 agents with the same snmpEngineID are compromised. SNMPv3 managers compliant with the SNMPv3 specification may not communicate with the non-compliant agents. The visible symptom of this problem may be an inexplicable loss of communication with those SNMPv3 agents.	Change the snmpEngineID of the affected SNMPv3 agents to unique values. Also, on those SNMPv3 agents, you will need to recreate new security keys from the security passwords. (This is an SNMP agent-specific process, so consult the documentation for your SNMP agent.) Finally, when finished, restart the SNMPv3 agent to activate the changes.

What Security Analyzer Detects	What This Means	Why It Matters	How To Fix
The SNMP Security Analyzer can detect if the snmpEngineBoots and snmpEngineTime values on an SNMPv3 agent go "backwards."	The SNMPv3 specification requires that each SNMPv3 agent maintain snmpEngineBoots and snmpEngineTime values, and that these values be continually increasing. If either of these values decrease, then the time values go "backwards." The time values are said to have "rolled back."	SNMPv3 security limits replay attacks, where SNMPv3 requests are recorded and re-sent later. The snmpEngineBoots and snmpEngineTime values protect against replay attacks. If an SNMPv3 agent's snmpEngineBoots and snmpEngineTime values go backwards, then the SNMPv3 agent becomes more vulnerable to replay attacks. If an SNMPv3-compliant manager receives a reply from an SNMPv3 agent whose time values have gone backwards, it will assume the reply is a replay attack. The manager will be unable to trust the authenticity of the reply and will discard it. The visible sympton of this problem may be an inexplicable loss of communication with that SNMPv3 agent.	Clock rollbacks occur for one of two reasons: either the agent's stable-storage configuration is changed "out from under" the running agent using a non-SNMP mechanism, or the SNMP agent is defective. To resolve a situation where a running agent's stable-storage configuration is changed "out from under" it, again change the agent's stable-storage values of snmpEngineBoots and snmpEngineTime. Set these to values larger than their current values, then restart the SNMPv3 agent. If the SNMPv3 agent is defective and incorrectly maintains the snmpEngineBoots and snmpEngineTime values, then the recommended fix is to upgrade the SNMPv3 agent to a fixed version. After deploying either of the fixes above, force the SNMP manager to rediscover the agent to reset the manager's notion of the agent's clock. If you are unable to upgrade the SNMPv3 agent, then you should evaluate your SNMP security policy to decide the best course of action. Possibilities include: • continuing to use the vulnerable SNMPv3 agent without authentication; • restricting the information accessible on that agent; • restricting the management stations the SNMP agent will accept requests from; • ignoring traps received from that agent because the type, content, and timeliness of the the alerts cannot be verified. (SNMP Inform messages can be used instead of traps.)

What Security Analyzer Detects	What This Means	Why It Matters	How To Fix
The SNMP Security Analyzer can detect if the snmpEngineBoots and snmpEngineTime values on an SNMPv3 agent have reached their maximum values and are not continually increasing.	The SNMPv3 specification requires that each SNMPv3 agent maintain snmpEngineBoots and snmpEngineTime values, and that these values be continually increasing. If either of these values have reached their maximum values, then they cannot keep increasing. The time values are said to be "latched" at their maximum values. Note that in practice, it is extremely unlikely the snmpEngineBoots and snmpEngineTime will ever reach their maximum values and "latch." The mere occurance of this condition is an anomaly that should be investigated.	SNMPv3 security limits replay attacks, where SNMPv3 requests are recorded and re-sent later. The snmpEngineBoots and snmpEngineTime values are used to limit replay attacks. If an SNMPv3 agent's snmpEngineBoots and snmpEngineTime values are latched at their maximum values, then the SNMPv3 agent becomes more vulnerable to replay attacks. If an SNMPv3-compliant manager receives a reply from an SNMPv3 agent whose time values have latched and are not increasing, it will assume the reply is a replay attack. The manager will be unable to trust the authenticity of the reply and will discard it. The visible sympton of this problem may be an inexplicable loss of communication with that SNMPv3 agent.	The part of the SNMPv3 standard that discusses recovery from this condition (RFC 3414, Section 2.2.2) says: "In order to reset an SNMP engine whose snmpEngineBoots value has reached the value 2147483647, manual intervention is required. The engine must be physically visited and re-configured, either with a new snmpEngineID value, or with new secret values for the authentication and privacy protocols of all users known to that SNMP engine. Note that even if an SNMP engine re-boots once a second that it would still take approximately 68 years before the max value of 2147483647 would be reached."

What Security Analyzer Detects	What This Means	Why It Matters	How To Fix
The SNMP Security Analyzer can detect SNMPv1 or SNMPv2c agents that use community strings that allow write access to MIB objects.	SNMPv1 or SNMPv2 can be used to change the state of network devices, including the possibility of enabling or disabling them.	SNMPv1 or SNMPv2c community strings are sent in clear text, unencrypted, and are used to identify an authority who is sending the SNMP set request. Because the clear-text community strings are insecure, they are suceptible to interlopers "snooping" the SNMP messages to obtain the identity of the sending authority. The interloper could then send malicious SNMPv1 or SNMPv2c requests to your agent, posing as that authority.	The recommended fix is to upgrade to an SNMPv3-capable agent with security features, enable and use SNMPv3 messaging with security features, and disable the use of SNMPv1 or SNMPv2c for write access. If upgrading your agent to support SNMPv3 and enabling security is not possible, then you should evaluate your SNMP security policy to decide the best course of action. Possibilities include: • continuing to use the vulnerable SNMPv1/SNMPv2c agent as-is; • restricting the management stations from which the SNMP agent will accept requests; • isolating the SNMPv1/SNMPv2c agent to its own protected subnet;
The SNMP Security Analyzer can detect SNMPv1 or SNMPv2c agents that use community strings that allow read access to MIB objects.	SNMPv1 or SNMPv2 can be used to observe the configuration of network devices, possibly observing sensitive information.	SNMPv1 or SNMPv2c community strings are sent in clear text, unencrypted, and are used to identify the authority who is sending the SNMP get request. Because the clear-text community strings are insecure, they are suceptible to interlopers "snooping" the SNMP messages to obtain the identity of the sending authority and the contents of the SNMP message. The interloper could then send malicious SNMPv1 or SNMPv2c requests to your agent, posing as that authority to request additional information.	The recommended fix is to upgrade to an SNMPv3-capable agent with security features, enable and use SNMPv3 messaging with security features, and disable the use of SNMPv1 or SNMPv2c for read access. If upgrading your agent to support SNMPv3 and enabling security is not possible, then you should evaluate your SNMP security policy to decide the best course of action. Possibilities include: • continuing to use the vulnerable SNMPv1/SNMPv2c agent as-is; • restricting the management stations from which the SNMP agent will accept requests; • isolating the SNMPv1/SNMPv2c agent to its own protected subnet;

What Security Analyzer Detects	What This Means	Why It Matters	How To Fix
The SNMP Security Analyzer can detect SNMPv3 agents that accept SNMPv3 write requests but do not require secure verification of the identity of the requestor (SNMPv3 with no authentication)	The SNMPv3 agent will respond to an SNMP write request without securely verifying the identity of the authority sending the request. An interloper could exploit this ability to change the state of the SNMP agent, including the possibility of enabling or disabling it.	The SNMPv3 specification requires strong authentication mechanisms for securely identifying the authority sending the SNMP request. If these are not used, the SNMP agent is vulnerable to write requests from interlopers who may send malicious write requests to your SNMP agent.	Enable and use strong authentication on your SNMPv3 agent. Configure the agent to require authenticated SNMPv3 messages for write (set) requests, or limit what agent information can be changed for SNMP requests having no security.
The SNMP Security Analyzer can detect SNMPv3 agents that accept SNMPv3 read requests but do not require secure verification of the identity of the requestor (SNMPv3 with no authentication)	The SNMPv3 agent will respond to an SNMP read request without securely verifying the identity of the authority sending the request. An interloper could exploit this ability to observe the state of the SNMP agent, including sensitive information.	The SNMPv3 specification requires cryptographic mechanisms for securely identifying the person sending the SNMP request. If these are not used, the SNMP agent is vulnerable to read requests from interlopers who may send malicious read requests to your SNMP agent.	Enable and use strong authentication on your SNMPv3 agent. Configure the agent to require authenticated SNMPv3 messages for read (get) requests, or limit what agent information can be observed for SNMP requests having no security.
The SNMP Security Analyzer collects information about security-related parameters on SNMPv3 agents in your networks.	SNMPv3 supports remote configuration of the SNMPv3 security. The security-related information is stored in SNMP tables in the agent. The Security Analyzer collects all the security-related information it can, and reports it.	SNMPv3 agents may contain stale or forgotten security-related information, such as SNMPv1/SNMPv2c community strings, SNMPv3 users, security groups, and MIB views. Reporting this information helps helps you find and remove these before they can be exploited by an interloper.	Edit the SNMP security-related tables to delete the extra information.