

[UNIX] Vulnerability in Encrypted Loop Device for Linux

Source: <http://www.derkeiler.com/Mailing-Lists/Securiteam/2002-01/0012.html>

From: support@securiteam.com

Date: 01/04/02

From: support@securiteam.com

To: list@securiteam.com

Date: Fri, 4 Jan 2002 10:48:12 +0100 (CET)

The following security advisory is sent to the securiteam mailing list, and can be found at the SecuriTeam web site: <http://www.securiteam.com>

-- promotion

When was the last time you checked your server's security?

How about a monthly report?

<http://www.AutomatedScanning.com> -- Know that you're safe.

Vulnerability in Encrypted Loop Device for Linux

SUMMARY

This text describes a security hole that was found in encrypted loop device for Linux. An attacker can modify the content of the encrypted device without being detected (see section 2). This text proposes to fix the hole by authenticating the device (see section 3).

DETAILS

Threat model (Section 1):

Encrypting a disk device aims to protect against an off-line attacker who would be able to access the disk between two legitimate mounts.

It is not against an attacker who has access to the running computer when the encrypted device is mounted as either (i) the attacker is root and it can access the encrypted device anyway or (ii) he is an unprivileged user and can be stopped with UNIX's right management (i.e. user/group).

Attack description (Section 2):

The vulnerability of encrypted loop device is due to its lack of authentication. The aim of encryption is to make the data unreadable for anybody who does not know the key. It does not prevent an attacker from

Securiteam: [UNIX] Vulnerability in Encrypted Loop Device for Linux

modifying the data. People assume that an attacker will not do it because the attacker would not be able to choose the resulting clear text. However, this section shows that the attacker can choose the resulting clear text to some extent and that modifying the cipher text data may be interesting even if the attacker ignores the result.

This attack is only applicable to device storing data that are reused across mounts: most file-system (e.g. ext2, reiserfs, ext3) but not swap. In some systems, encrypted devices are stored in the same location than the encrypted disk containing the operating system. For those systems the attacker who can access the encrypted device, can easily modify the OS to gain access (e.g. kernel) independently of the encrypted device.

To insert random data (Section 2.1):

If the attacker modifies the cipher text without choosing the resulting clear text, it will likely produce random data. The legitimate user will not detect the modification and will use them as if they were valid. As they likely appear random, it will result of a Denial of Service.

To insert chosen data (Section 2.2):

The encryption mode used by encrypted loop device is CBC [oST81,sec 5.3]. CBC allows cut/past attacks i.e. the attacker can cut encrypted data from one part of the device and paste them in another location. As both data sections have been encrypted by the same key, the clear text will not be completely random data.

This lack of authentication is not a CBC flaw. Authentication is not considered an aim of the encryption mode, so most modes (e.g. ECB, CFB, and OFB) do not authenticate the data. To use another mode would be flawed in the same way except if they explicitly protect against forgery. Recently some modes including authentication popped up to speed up the encryption / authentication couple but as far as we know, they are all patented.

In very short, encrypting with CBC is $C_n = \text{Enc}(C_{n-1} \text{ xor } P_n)$ where $\text{Enc}(x)$ is encrypting x , P_n is the n th block of plain text and C_n the n th block of cipher text. For the first block, C_{n-1} is an Initial vector (a.k.a IV) which may be public and must be unique for a given key. The decryption is $P_n = \text{Dec}(C_n) \text{ xor } C_{n-1}$. See [oST81, sec 5.3] for a longer description of CBC.

If the attacker copies s blocks from the location m to n (a.k.a $[C_n, \dots, C_{n+s-1}] == [C_m, \dots, C_{m+s-1}]$), P_{n+1} up to P_{n+s-1} will be the same as P_{m+1} to P_{m+s-1} and P_n will likely appear random. C_n (i.e. C_m) will be decrypted as $P_n = \text{Dec}(C_m) \text{ xor } C_{n-1}$ but C_{n-1} and C_{m-1} are different so P_n will likely appear random. Nevertheless $P_{n+1} = \text{Dec}(C_{n+1}) \text{ xor } C_n = \text{Dec}(C_{m+1}) \text{ xor } C_m = P_{m+1}$, so $P_{n+1} = P_{m+1}$. Therefore, if the attacker has an idea of the content of a group of blocks in the device, he can copy them to the n th block, thus it can choose the content of it without being detected.

Securiteam: [UNIX] Vulnerability in Encrypted Loop Device for Linux

As a file-system is not designed to appear random, its content may be predictable to some extents (e.g. common directories and files, inode, superblock). The attacker may use such information to guess the contents and do a knowledgeable cut/past. For example, an attacker knowing the location of a password file may replace a password by another one that is already known.

Proposed fixes (Section 3):

We propose 2 types of fixes: one that authenticates at mount time (see section 3.1) and the other that authenticates at the cluster level (see section 3.2). The choice between the two (see section 3.4) is a user matter as it mostly depends on the access pattern on the encrypted device.

In the proposed fixes, the authentication is a MAC computed over the encrypted device. The MAC is HMAC[KBC97] combined with a configured hash function, preferably a well studied one such as SHA1[oST95] or MD5[Riv92]. The MAC secret key is derived from the pass-phrase via PKCS-5 key derivations ([Kal00,sec 5.1]).

Authenticating at mount time (Section 3.1):

As we need to authenticate the device across mounts and not while it is mounted (see section 1), it is sufficient to authenticate the whole device during mount operations. It slows down mount operations but they are rather infrequent so we consider the trade-off delay/security acceptable. The MAC is verified during mount operations and generated during unmount operations. It is not supposed to be valid while the device is mounted.

The MAC generation is done when unmounting the device. The MAC is computed over all the sectors of the device and the result is appended in the device file after all the sectors.

The MAC verification is done when mounting the device. The MAC is computed over all the sectors of the device. If the result is equal to the MAC appended to the block device, the verification succeed, else it failed.

The verification may fail (i) if an attacker attempted to modify the device during 2 legitimates mounts or (ii) if the device has not been cleanly unmounted (e.g. computer crash). It is impossible to automatically distinguish both cases with certainty. Therefore, if the verification fails, the user is notified and the mount operation may be stopped depending on configuration.

Authenticating at cluster level (Section 3.2):

To authenticate the whole device at mount time, may be considered prohibitive by some users, so this section describes an alternative that authenticates the device at the cluster level. A cluster is a group of one or more sectors, the exact number depends on configuration. In this case, the MAC is verified each time a cluster is read from the disk and generated at each write.

Securiteam: [UNIX] Vulnerability in Encrypted Loop Device for Linux

If the device is not cleanly unmounted, the authentication of one or more cluster may fail (e.g. the super block). This case will be detected at mount time. Nevertheless, if an attacker forges data in the device, it will be detected only when the user read the modified data. The kernel will read the forged cluster and the authentication will fail. It may report it with a printk with a rate limiter, it is not clean, but we do not see any better way.

MAC location (Section 3.3):

Currently the encrypted loop file-system is stored in a regular file of a hosting file-system. Its size is a multiple of a sector size (i.e. 512 byte). The MAC could be stored in a separate file or included in the regular file. To store the MAC in a separate file, generates problems while managing the loop device file (e.g. copy, backup). The administrator must not forget to copy the MAC file when he copies the device file, else the copied device will not be usable anymore. To store the MAC in the same file as the clusters does not have this disadvantage.

Comparison (Section 3.4):

To authenticate at the cluster level will increase the access time of each cluster but will not affect mount operation. The exact increase depends on the MAC and encryption algorithms. As a rule of thumb, MAC algorithms are typically 3 times faster than encryption ones so the time dedicated to cryptography for each block will increase by around 30%. To authenticate at mount time will largely slow down the mount operations but will not affect every access once mounted.

The authentication at mount time will detect forgery at mount time, whereas the alternative detects it only when the forged cluster is read, possibly a long time after the modification. Users may consider that it is easier to diagnose who forged it if they have a better idea of when the attack occurred.

To authenticate the whole device at mount time requires a single MAC per device, so the space overhead (typically 16 byte) is negligible compared to the device's size. To authenticate at the cluster level requires a MAC per cluster, it is significantly more, but some people may consider it still negligible, especially with cheap disks.

The choice between the two mostly depends on the access pattern on the encrypted device. If the device is used for interactive purpose, the increased access latency may be unsuitable. If the access latency is important or if every block is frequently modified, to authenticate only once at mount time may be more interesting. If the user cannot stand long mount operations, to authenticate at cluster level will be more suitable.

As only the final user knows the type access made on his encrypted device, he should be the one able to choose between the two.

Acknowledgments (Section 4):

Thanks to Andy Kleen and Phil Swan for their useful comments.

Securiteam: [UNIX] Vulnerability in Encrypted Loop Device for Linux

Conclusion (Section 5):

This text described a vulnerability in encrypted loop device that allows an attacker to modify the encrypted device without being detected (see section 2). We propose a fix, which authenticate the whole device during mount operation (see section 3.1). This fix slows down mount operations but we consider the trade-off longer delay vs. additional security very reasonable, as mount operations are rather infrequent. We propose another fix that authenticate at cluster level for people who cannot stand long mount operation. The choice between the two is a final user matter.

The authentication may be optionally disabled thus if a user considers the trade-off delay/security not in favor of security, he may choose to be vulnerable to this attack and disable it. Nevertheless, the author thinks encrypted loop device must be secure by default.

ADDITIONAL INFORMATION

The information has been provided by <<mailto:jme@off.net>> Jerome Etienne.

References:

[Kal00]

B. Kaliski. Pkcs 5: Password-based cryptography specification version 2.0. Request For Comment (Informational) RFC2898, September 2000.

[KBC97]

H. Krawczyk, M. Bellare, and R. Canetti. Hmac: Keyed-hashing for message authentication. Request For Comment (Informational) RFC2104, February 1997.

[oST81]

National Institute of Standards and Technology. implementing and using the nbs data encryption standard. Federal information processing standards fips74, April 1981.

[oST95]

National Institute of Standards and Technology. Secure hash standard. Federal information processing standards fips180-1, April 1995.

[Riv92]

R. Rivest. The md5 message-digest algorithm. Request For Comment (Informational) RFC1321, April 1992.

=====

This bulletin is sent to members of the SecuriTeam mailing list.
To unsubscribe from the list, send mail with an empty subject line and body to:

Securiteam: [UNIX] Vulnerability in Encrypted Loop Device for Linux

list-unsubscribe@securiteam.com

In order to subscribe to the mailing list, simply forward this email to: list-subscribe@securiteam.com

=====
=====

DISCLAIMER:

The information in this bulletin is provided "AS IS" without warranty of any kind.

In no event shall we be liable for any damages whatsoever including direct, indirect, incidental, consequential, loss of business profits or special damages.

- *Previous message:* support@securiteam.com: "[NEWS] USENIX, Security 2002"
- *Messages sorted by:* [\[date \]](#) [\[thread \]](#) [\[subject \]](#) [\[author \]](#) [\[attachment \]](#)