



# Tumbleweed Email Firewall Remote Stack Overflow

04-July-2006

## Summary

Tumbleweed's Email Firewall (EMF) blocks spam and viruses, phishing and email fraud, and keeps hackers from compromising your network. To ensure compliance with government and industry regulations like HIPAA, GLBA, Sarbanes-Oxley and Safe Harbor (EU), MailGate Email Firewall provides sophisticated filtering, monitoring, encryption and reporting capabilities. According to product literature, Tumbleweed is used by: over 150 healthcare providers, the Department of Defense, the Department of Homeland Security, all four branches of the US Military, state and local governments internationally, 8 of the top 10 US banks, 4 of the top 5 Canadian banks, and 6 of the top 10 European banks.

Tumbleweed's EMF Decomposer, a component that decompresses incoming e-mail attachments, has three separate vulnerabilities within its LHA processing routines. The first issue causes the LHA processing engine to exhibit a stack-buffer overflow while processing extended-header filenames. The second issue is a stack overflow while processing LHA extended-header directory names. The third issue is a buffer overflow during a sprintf call while processing long filenames contained in an LHA archive.

## Impact

These vulnerabilities are present by default in Tumbleweed's Email Firewall. To exploit these vulnerabilities, an attacker only needs to send an e-mail to an organization running Tumbleweed; it is not necessary that the e-mail is opened. Successful exploitation of these vulnerabilities results in remote code execution with the full privileges of the MMSDecompose process. The default settings allow an attacker to obtain super-user access to the machine. Since these vulnerabilities are stack based overflows, exploits can be made to work reliably.

## Affected software

Tumbleweed Email Firewall (All Versions)

## Credit

These vulnerabilities were researched by Ryan Smith.

## Contact

advisories@hustlelabs.com




## Details

In the following code segment, the program reads a word-sized value from the file: the LHA extended-header size. If the LHA header level is equal-to 1 then the program will read in more data to a buffer. Next, depending on the value of the extended-header-type byte, the program branches to an area of code that handles the specific type or a generic handler if it's an unrecognized type.

```

01902488 DIGEST_EXTENDED_HEADERS:                ; CODE XREF: [
01902488                                     ; DecomposeLZF
01902488      call    GetWORD
01902488      movzx  esi, ax
01902488      test   esi, esi
01902488      jz     loc_1902653
01902488      cmp   [ebx+LzHdr.level], 2
01902488      jz     short loc_19024FD
01902488      lea   eax, [esp+141Ch]
01902488      sub   eax, glb_CurrFileDataPtr
01902488      cmp   eax, esi
01902488      jl     dissl
01902488      push  edi                ; FILE *
01902488      mov   eax, glb_CurrFileDataPtr
01902488      push  esi                ; size_t
01902488      push  1                  ; size_t
01902488      push  eax                ; void*
01902488      call  _fread
01902488      add   esp, 10h
01902488      cmp   eax, esi
01902488      jb     dissl
01902488
01902488      loc_19024FD:                ; CODE XREF: [
01902488      mov   ecx, glb_CurrFileDataPtr
01902488      xor   eax, eax
01902488      inc   glb_CurrFileDataPtr
01902488      mov   al, [ecx]
01902488      cmp   eax, 54h
01902488      ja     short ext_unknown
01902488      xor   ecx, ecx
01902488      mov   cl, ds:byte_190281C[eax]
01902488      jmp   ds:extended_hdr_switch[ecx*4]

```



If the extended-header-type byte is equal to 0x01, then the following code parses the data for the header. The size allocated for this buffer is equal-to 0x100 bytes, but there is no length restriction. Thus, an attacker can supply a value greater-than 0x100 in an archive file to cause a buffer-overflow.

```

01902540 ext_filename_hdr:                ; CODE XRE
01902540                                     ; DATA XRE
01902540      xor   edx, edx
01902540      lea   ecx, [esi-3]
01902540      test  ecx, ecx
01902540      jle   short loc_1902556
01902540
01902540      loc_1902540:                ; CODE XRE
01902540      mov   eax, glb_CurrFileDataPtr
01902540      inc   edx
01902540      mov   al, [eax]
01902540      mov   [edx+ebx+13h], al
01902540      inc   glb_CurrFileDataPtr
01902540      cmp   edx, ecx
01902540      jl    short loc_1902540
01902540
01902540      loc_1902556:                ; CODE XRE
01902540      mov   byte ptr [ebx+esi+11h], 0
01902540      jmp   DIGEST_EXTENDED_HEADERS

```



As well, if the extended-header-type byte is equal to 0x02, then this next piece of code parses the data for the header. The size allocated for this buffer is equal-to 0x100 bytes; however, there is no restriction for the number of bytes the program will copy from the file, to the buffer. An attacker can supply a value greater-than 0x100 to cause a buffer-overflow.

```

01902574 ext_directory_hdr:                ; CODE XREF: DecomposeL
01902575                                     ; DATA XREF: .text:0190
01902576                                     ; CODE XREF: DecomposeL
01902577 mov     ecx, ecx
01902578 lea   ebp, [esi-3]
01902579 test  ebp, ebp
0190257A jle    short loc_190257F
0190257B
0190257C loc_1902569:                       ; CODE XREF: DecomposeL
0190257D mov     eax, glb_CurrFileDataPtr
0190257E inc     ecx
0190257F mov     al, [eax]
01902580 mov     byte ptr [esp+ecx+141Ch+ucCkSum+3], al
01902581 inc     glb_CurrFileDataPtr
01902582 cmp     ecx, ebp
01902583 jl     short loc_1902569
01902584
01902585 loc_190257F:                       ; CODE XREF: DecomposeL
01902586 lea   eax, [esp+141Ch+var_1400]
01902587 push  eax
01902588 mov     byte ptr [esp+esi+1420h+ucCkSum+1], 0
01902589 push  eax
0190258A call   sub_1908060
0190258B add     esp, 8
0190258C jmp    DIGEST_EXTENDED_HEADERS
01902593

```



The next code excerpt is responsible for concatenating a temporary directory and the filename strings, in order to derive the path for the output of decompression. The first variable to the sprintf call is the format string "%s/%s", the second variable is the temporary pathname, and the third is the filename. Of these parameters, the filename parameter is a user controlled value whose size should be no larger than 0x100 bytes (Including a null-termination character), according to the program's data structures. Ignoring the fact that the size restriction is not enforced, a temporary pathname greater-than 2 characters in length will allow a buffer overflow to occur even for archive names that are legitimately sized.

```

01905C74 cmp     dword_1913384, 0
01905C75 jz     short loc_1905C9D
01905C76 push  edi
01905C77 mov     eax, dword_1913384
01905C78 lea   ecx, [esp+118h+var_104]
01905C79 push  eax
01905C7A push  offset aSS ; "%s/%s"
01905C7B push  ecx ; char *
01905C7C call   _sprintf
01905C93

```






## Remediation

The code should be modified such that there is a standard maximum length of a path. If the path exceeds the maximum length, then the file should be rejected, or the path truncated.

Though Tumbleweed won't release a patch, they officially recommend the following actions:

- Stopping the EMF services
- Removing or renaming the wlha32.dll file, found in the EMF install directory.
- Restarting the EMF services



## Timeline of Events

04-July-2006 – Advisory draft

11-July-2006 – Vendor notification

24-July-2006 – Vendor released customer notification and the workaround



## Attributions

Images of Billy the Kid, Jesse James, Butch Cassidy, and Doc Holiday were taken from Wikipedia.

(<http://www.wikipedia.org>)

Image of the tumbleweed along a road was photographed by AV Smith, from the Galveston Arts Center. (<http://www.galvestonartscenter.org>)

Code and cross-reference screenshots captured using IDA

(<http://www.datarescue.com>).

Flawed code and marketing information obtained from Tumbleweed

(<http://www.tumbleweed.com>).

The Creative Commons license-notification image borrowed from

<http://www.creativecommons.org>.

## License

This work is licensed under the Creative Commons Attribution 2.5 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by/2.5/> or send a letter to Creative Commons, 543 Howard Street, 5th Floor, San Francisco, California, 94105, USA.

Attribution should be provided both in the form of a link or reference to <http://www.hustlelabs.com> and a copy of the researchers' names listed under the *Credit* section of this document.

All other trademarks and copyrights referenced in this document are the property of their respective owners.

