Playing Hide-and-Seek with Hash-DoS

Pascal Junod // HEIG-VD

Insomni'hack 2013, Geneva (Switzerland)







Overview

- DoS and Complexity attacks
- Hash DoS
- BTRFS
- Apache





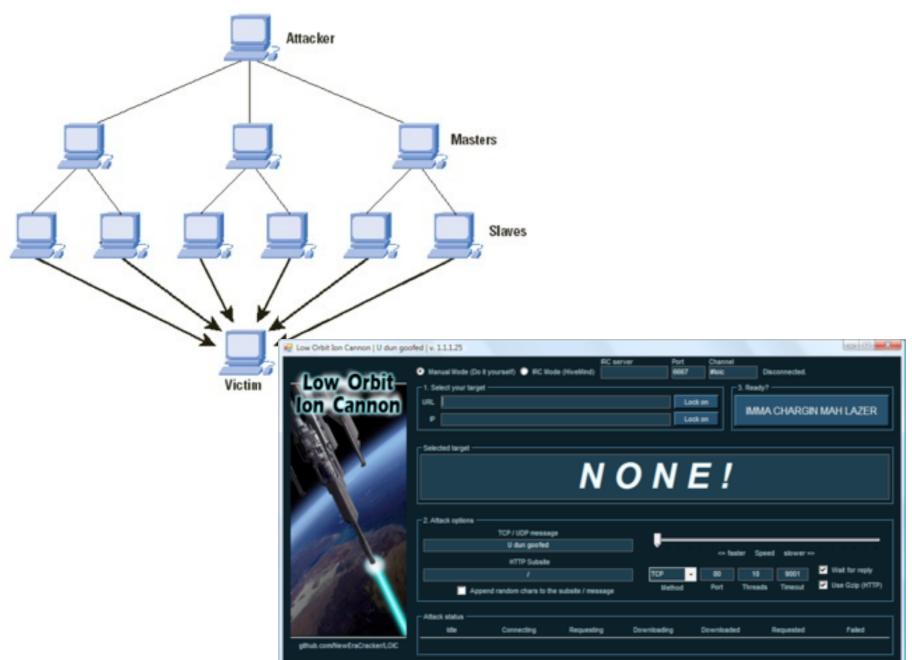
Denial of Service

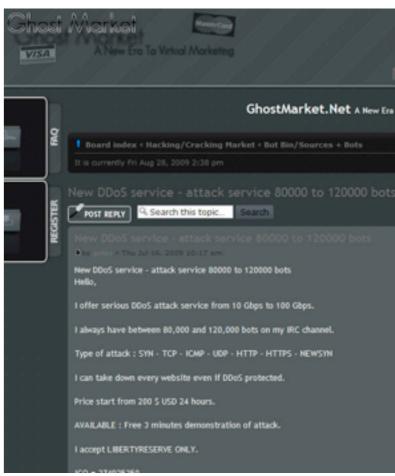
- Goal is to make a resource unavailable to its intended users.
- Many different ways to do it:
 - (Bad programmers)
 - DDoS with TCP SYN or HTTP queries flood
 - Exploitation of amplification mechanisms
 - Complexity attacks
 - ...





DDoS with TCP SYN or HTTP flood





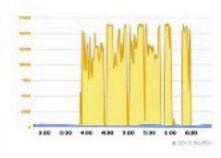




".ch" survives DDoS attack unscathed

January 10, 2013 / Roland Eugster

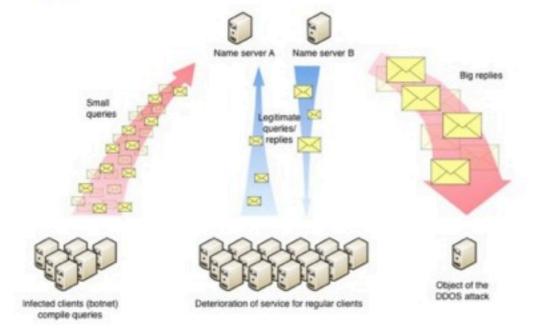
Since this morning (Thursday) all the Swiss name servers have been systematically abused in a bid to stop other websites from operating. "Distributed Denial of Service (DDoS)" attacks are nothing new, but this is the first time that the .ch infrastructure has been abused. Thanks to the high-quality operation of Switzerland's Internet by the SWITCH Foundation and SWITCH's rapid intervention, all .ch websites have remained accessible all the time.



Since 04:00 this morning, all the ".ch" name servers have been attacked by meaningless queries with an intensity that is many times that of the normal network load. The .ch zone has not been the object of the attack but just the means to the end. In abusing the Swiss name servers, the attackers are attempting to prevent various websites in the USA from operating and are setting out to cause damage their operators.

Stable operation guaranteed

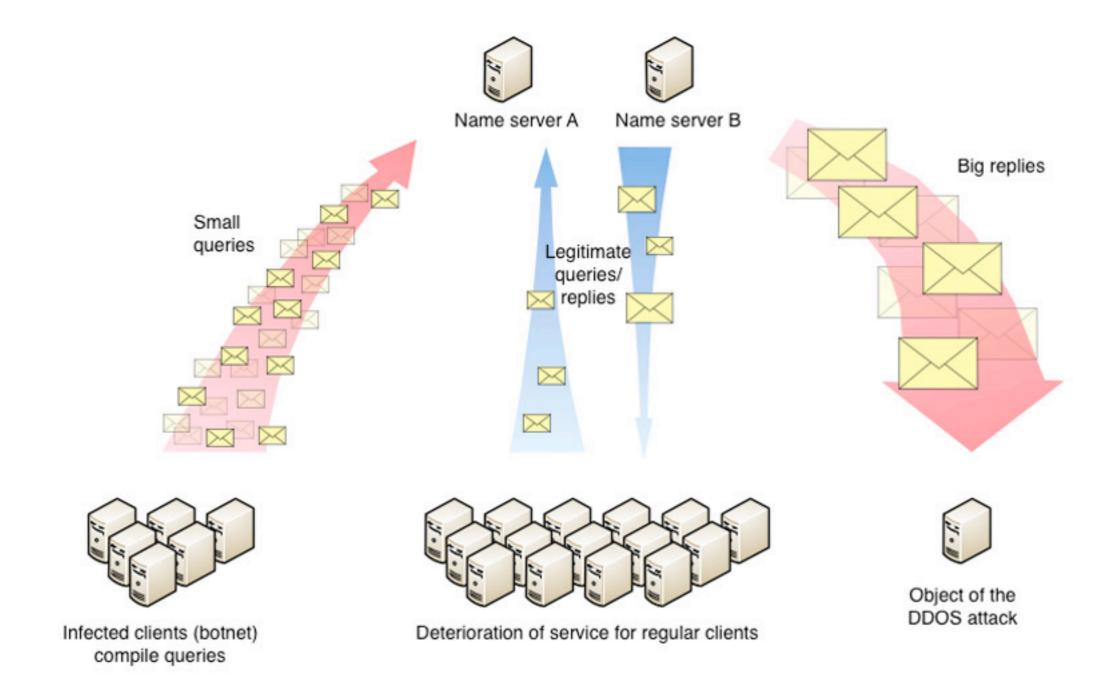
The attack – a standard "Distributed Denial of Service" attack – could have had far-reaching effects had there not been sufficient security precautions in place: if all the name servers are blocked, then no .ch websites can be accessed. Thanks to the rapid intervention of SWITCH's security team, it proved possible to defuse the situation. "We were prepared for such an emergency and were able to activate the necessary filters straight away and thus block the malicious traffic", explains Daniel Stirnimann who is responsible for the name server infrastructure. Since then, the load has been running at the normal level again, even though the attack is still ongoing.



Source: http://www.switch.ch/about/news/2013/ddos.html

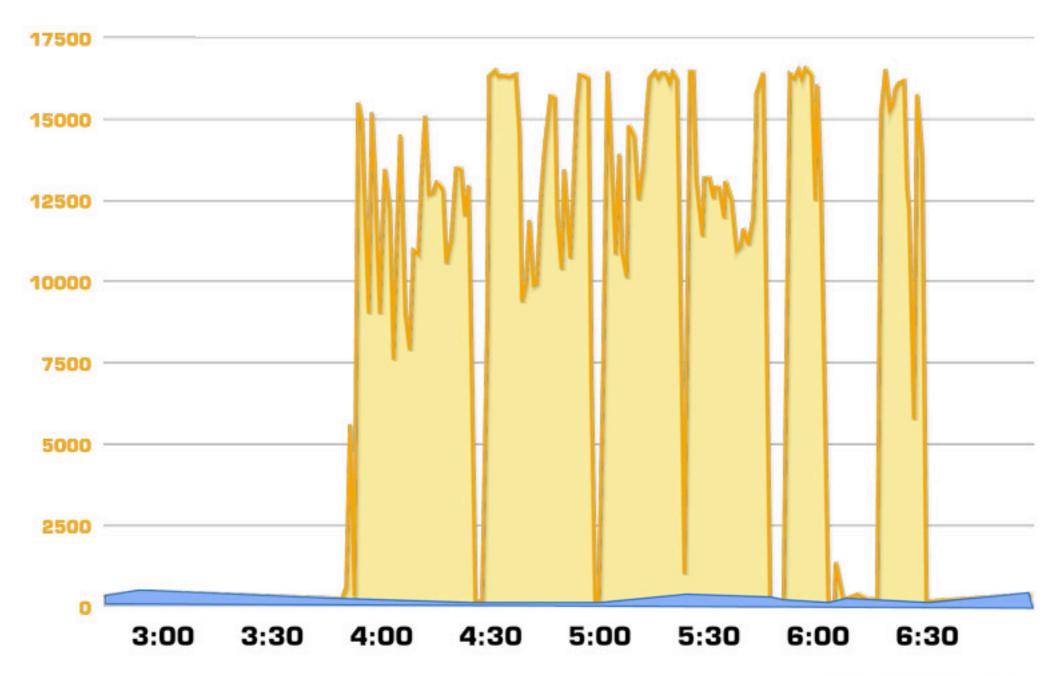






Source: http://www.switch.ch/about/news/2013/ddos.html





© 2013 SWITCH





----[Introduction

The purpose of this article is to show potential problems with intrusion detection systems (IDS), concentrating on one simple attack: port scans.

This lets me cover all components of such a simplified IDS. Also, unlike the great SNI paper (http://www.secnet.com/papers/IDS.PS), this article is not limited to network-based tools. In fact, the simple and hopefully reliable example port scan detection tool ("scanlogd") that you'll find at the end is host-based.



----[Data Structures and Algorithm Choice

When choosing a sorting or data lookup algorithm to be used for a normal application, people are usually optimizing the typical case. However, for IDS the worst case scenario should always be considered: an attacker can supply our IDS with whatever data she likes. If the IDS is fail-open, she would then be able to bypass it, and if it's fail-close, she could cause a DoS for the entire protected system.

Let me illustrate this by an example. In scanlogd, I'm using a hash table to lookup source addresses. This works very well for the typical case as long as the hash table is large enough (since the number of addresses we keep is limited anyway). The average lookup time is better than that of a binary search. However, an attacker can choose her addresses (most likely spoofed) to cause hash collisions, effectively replacing the hash table lookup with a linear search. Depending on how many entries we keep, this might make scanlogd not be able to pick new packets up in time. This will also always take more CPU time from other processes in a host-based IDS like scanlogd.

I've solved this problem by limiting the number of hash collisions, and

Let me illustrate this by an example. In scanlogd, I'm using a hash table to lookup source addresses. This works very well for the typical case as long as the hash table is large enough (since the number of addresses we keep is limited anyway). The average lookup time is better than that of a binary search. However, an attacker can choose her addresses (most likely spoofed) to cause hash collisions, effectively replacing the hash table lookup with a linear search. Depending on how many entries we keep, this might make scanlogd not be able to pick new packets up in time. This will also always take more CPU time from other processes in a host-based IDS like scanlogd.

more research might be needed.





Denial of Service via Algorithmic Complexity Attacks

Scott A. Crosby scrosby@cs.rice.edu Dan S. Wallach dwallach@cs.rice.edu

Department of Computer Science, Rice University

Senio W

We present a new class of low-bandwidth denial of service attacks that exploit algorithmic deficiencies in many common applications' data structures. Frequently used data structures have "average-case" expected running time that's far more efficient than the worst case. For example, both binary trees and hash tables can degenerate to linked lists with carefully chosen input. We show how an attacker can effectively compute such input, and we demonstrate attacks against the hash table implementations in two versions of Perl, the Squid web proxy, and the Bro intrusion detection system. Using bandwidth less than a typical dialup modem, we can bring a dedicated Bro server to its knees; after six minutes of carefully chosen packets, our Bro server was dropping as much as 71% of its traffic and consuming all of its CPU. We show how modern universal hashing techniques can yield performance comparable to commonplace hash functions while being provably secure against these attacks.

sume O(n) time to insert n elements. However, if each element hashes to the same bucket, the hash table will also degenerate to a linked list, and it will take $O(n^2)$ time to insert n elements.

While balanced tree algorithms, such as red-black trees [11], AVL trees [1], and treaps [17] can avoid predictable input which causes worst-case behavior, and universal hash functions [5] can be used to make hash functions that are not predictable by an attacker, many common applications use simpler algorithms. If an attacker can control and predict the inputs being used by these algorithms, then the attacker may be able to induce the worst-case execution time, effectively causing a denial-of-service (DoS) attack.

Such algorithmic DoS attacks have much in common with other low-bandwidth DoS attacks, such as stack smashing [2] or the ping-of-death ¹, wherein a relatively short message causes an Internet server to crash or misbehave. While a variety of techniques can be used to address these DoS attacks, common industrial practice still allows bugs like these



Attacks on Web Application
Platforms





Alexander "alech" Klink n.runs AG Julian "zeri" Wälde TU Darmstadt



#hashDoS

December 28th, 2011. 28th Chaos Communication Congress. Berlin, Germany.





Hash-flooding DoS reloaded: attacks and defenses

Jean-Philippe Aumasson, Kudelski Group

Daniel J. Bernstein, University of Illinois at Chicago

Martin Boßlet, freelancer



Hash-flooding DoS reloaded: attacks and defenses

Jean-Philippe Aumasson, Kudelski Security (NAGRA)

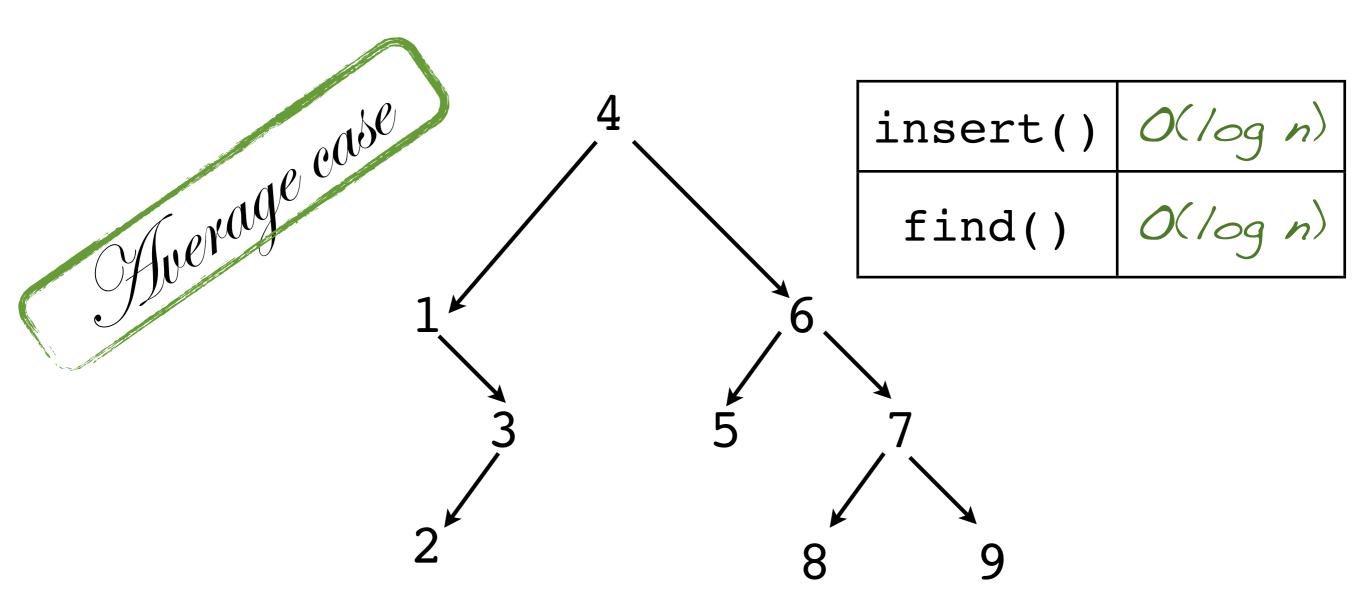
D. J. Bernstein, University of Illinois at Chicago & Technische Universiteit Eindhoven

Martin Boßlet, Ruby Core Team



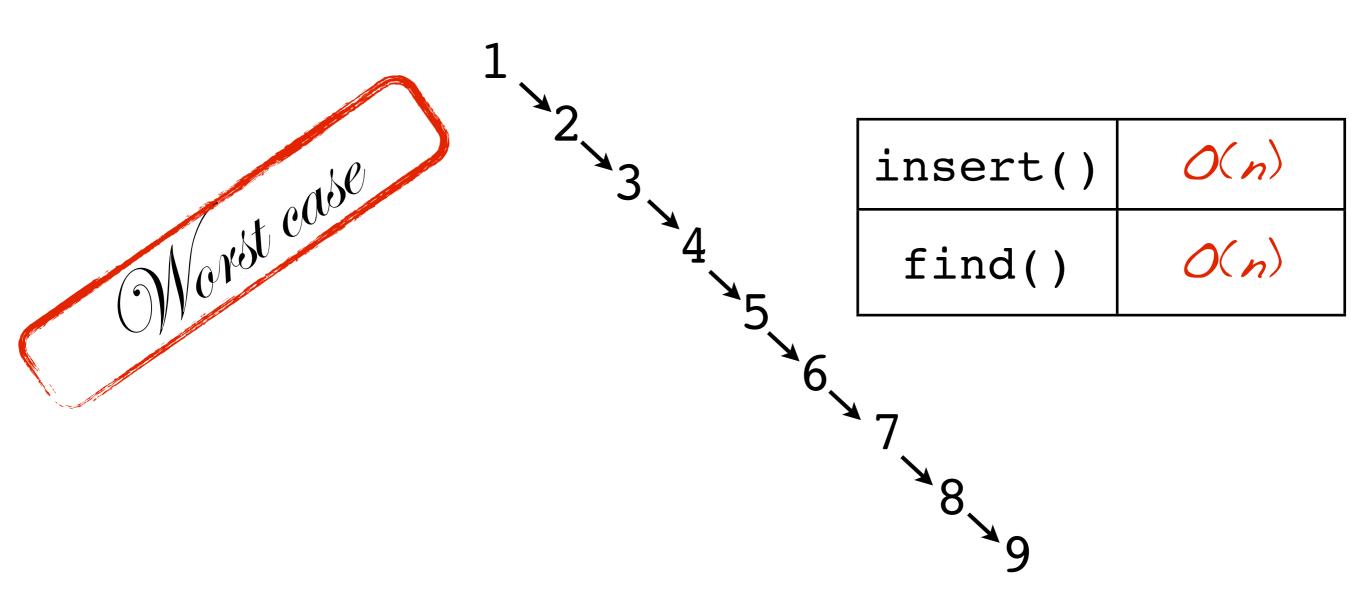


Example: Binary Tree



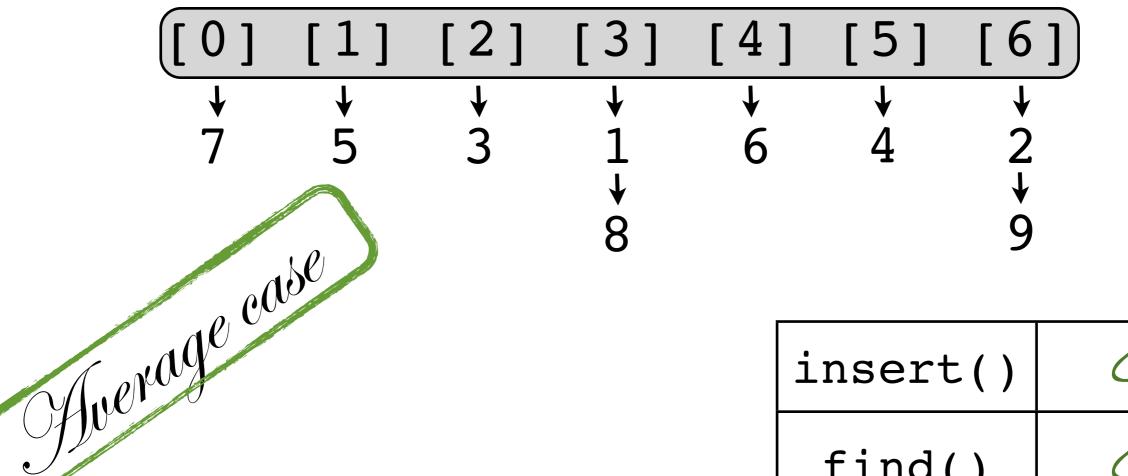


Example: Binary Tree





Example: Hash Table



uint8_t hash (uint32_t e) return (e*3) % 7;

insert()	0(1)
find()	0(1)



Example: Hash Table

```
[6]
                   [2]
                                      [5]
                          [3]
monst case
                                                   O(n)
                                     insert()
                                 34
                                                   O(n)
       uint8_t hash (uint32_t e)
                                      find()
         return (e*3) % 7;
```

Pascal Junod, «Playing Hide-and-Seek with Hash-DoS»

nsormi'hack 2013, March 22nd, 2013, Geneva (Switzerland)





Hash DoS





Multi-collisions

- Goal of a Hash-DoS: trigger the worstcase!
- In order to do it, one must find **multi-collisions** on the hash function, i.e., inputs $x_1, x_2, x_3, ..., x_n$ such that $H(x_1) = H(x_2) = H(x_3) = ... H(x_n)$
- Most simple hash functions allow this easily.





Randomization

- A possible counter-measure consists in randomizing the hash function, keeping secret the random parameter.
- Not all hash functions and constructions are OK...
- For instance, look at h := H(x)+r, where r is the random value and where multicollisions on H(.) are easy to find.





Better solution

- Use a data structure that has good worst case complexities, like red-black trees.
- Performance loss vs. security loss

Red-black tree

From Wikipedia, the free encyclopedia



This article needs additional citations for verification. Please help improve this article by adding citations to reliable sources. Unsourced material may be challenged and removed. (October 2012)

A red-black tree is a type of self-balancing binary search tree, a data structure used in computer science.

The self-balancing is provided by painting each node with one of two colors (these are typically called 'red' and 'black', hence the name of the trees) in such a way that the resulting painted tree satisfies certain properties that don't allow it to become significantly unbalanced. When the tree is modified, the new tree is subsequently rearranged and repainted to restore the coloring properties. The properties are designed in such a way that this rearranging and recoloring can be performed efficiently.

The balancing of the tree is not perfect but it is good enough to allow it to guarantee searching in $O(\log n)$ time, where n is the total number of elements in the tree. The insertion, and deletion operations, along with the tree rearrangement and recoloring are also performed in $O(\log n)$ time.^[1]

Red-black tree Tree Type Invented 1972 Invented by Rudolf Bayer Time complexity in big O notation Average Worst case O(n) O(n) Space O(log n) Search O(log n) Insert O(log n) O(log n) Delete O(log n) O(log n)

Red-black tree Type Tree Invented 1972 Rudolf Bayer Invented by Time complexity in big O notation Worst case Average. O(n) Space O(n) Search O(log n) O(log n) O(log n) O(log n) Insert O(log n) O(log n) Delete

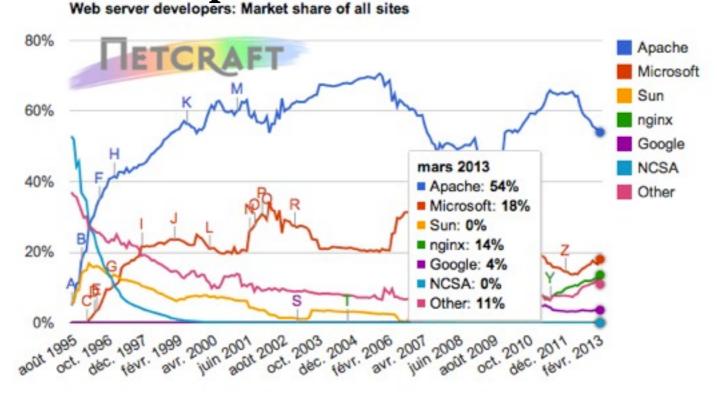
Tracking the color of each node requires only 1 bit of information per node because there are only two colors. The tree does not contain any other data specific to its being a red-black tree so its memory footprint is almost identical to classic (uncolored) binary search tree. In many cases the additional bit of information can be stored at no additional memory cost.





Example: nginx

- HTTP server / reverse proxy
- Written by Igor Sysoev since 2002
- High performances, low CPU/memory footprint compared to competitors







Examples: nginx

- All data coming from configuration files are stored in hash tables
 - Very efficient, but not secure

```
#define ngx_hash(key, c) ((ngx_uint_t) key * 31 + c)
```

- All data coming from outside are stored in red-black trees.
 - A bit less efficient, but mastered worstcase









- According to Wikipedia, the «B-tree file system» is a GPL-licensed experimental file-system for Linux.
- Several interesting features à la ZFS



Features

As of Linux 3.6 (released 30 September 2012), Btrfs implements: [22][23]

- Online defragmentation
- Online volume growth and shrinking
- Online block device addition and removal
- Online balancing (movement of objects between block devices to balance load)
- Offline filesystem check
- Online data scrubbing for finding errors and automatically fixing them for files with redundant copies
- RAID0, RAID1, and RAID10
- Subvolumes (one or more separately mountable filesystem roots within each physical partition)
- Transparent compression (zlib and LZO)
- Snapshots (read-only^[24] or copy-on-write clones of subvolumes)
- File cloning (copy-on-write on individual files, or byte ranges thereof)
- Checksums on data and metadata (CRC-32C^[25])
- In-place conversion (with rollback) from ext3/4 to Btrfs^[26]
- File system seeding^[27] (Btrfs on read-only storage used as a copy-on-write backing for a writeable Btrfs)
- . Block discard support (reclaims space on some virtualized setups and improves wear leveling on SSDs with TRIM)
- Send/receive (saving diffs between snapshots to a binary stream)^[28]
- Hierarchical per-subvolume quotas^[29]

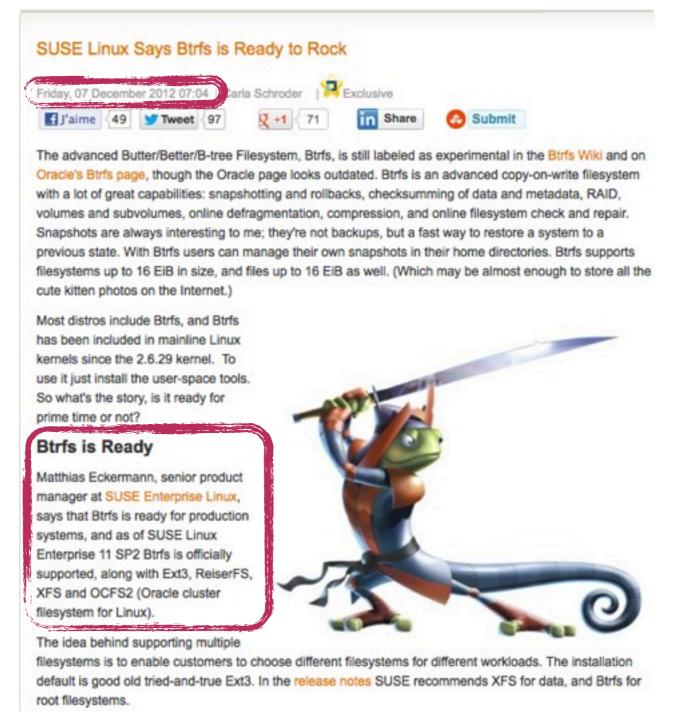
Planned features include:

- Online filesystem check^[30]
- Very fast offline filesystem check[citation needed]
- Parity-based RAID (RAID5 and RAID6)^[30]
- Object-level RAID0, RAID1, and RAID10^[citation needed]
- Incremental dumps [citation needed]
- Ability to handle swap files and swap partitions^[31]
- Data deduplication^{[4][32]}
- Encryption^{[4][31]}

Source: http://en.wikipedia.org/wiki/Btrfs#Features







Source: http://www.linux.com/news/enterprise/systems-management/677226-suse-linux-says-btrfs-is-ready-to-rock



BTRFS Directory Entry

Directories

Directories are indexed in two different ways. For filename lookup, there is an index comprised of keys:

Directory Objectid BTRFS_DIR_ITEM_KEY 64 bit filename hash

The default directory hash used is crc32c, although other hashes may be added later on. A flags field in the super block will indicate which hash is used for a given FS.

The second directory index is used by readdir to return data in inode number order. This more closely resembles the order of blocks on disk and generally provides better performance for reading data in bulk (backups, copies, etc). Also, it allows fast checking that a given inode is linked into a directory when verifying inode link counts. This index uses an additional set of keys:

Directory Objectid BTRFS_DIR_INDEX_KEY Inode Sequence number

The inode sequence number comes from the directory. It is increased each time a new file or directory is added.

Source: https://btrfs.wiki.kernel.org/index.php/Btrfs_design





<math-nerd-session>





CRC32c

- Cyclic Redundancy Check (CRC)
 - Linear error-detecting code
 - Non-cryptographic!
 - (Multi-)collisions are easy to find
 - Pre-images are easy to find





CRCs

- ▶ Easily said: CRC(x) is the remainder of the division of x interpreted as a polynomial over GF(2) by the generator polynomial.
- Example:
 - $x = 0 \times 77$, $G(x) = x^4 + x + 1$





CRCs

- 0x77 <=> 0b01110111

Generator polynomial

Quotient, that we forget

CRC4(0x77) = 0xE





CRCs & Multi-Collisions

$$x^{6}+x^{5}+x^{4}+x^{2}+x+1 = (x^{4}+x+1)(x^{2}+x+1)+x^{3}+x^{2}+x$$

$$x^{4}+x^{3}+x^{2}+1 = (x^{4}+x+1)1+x^{3}+x^{2}+x$$

$$x^{5}+x^{3} = (x^{4}+x+1)x+x^{3}+x^{2}+x$$

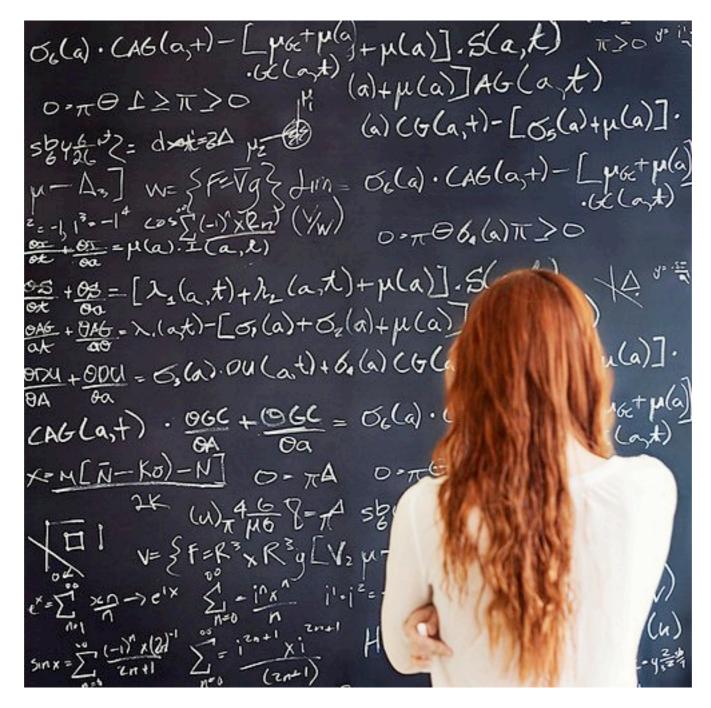
$$x^{5}+x^{4}+x^{3}+x+1 = (x^{4}+x+1)(x+1)+x^{3}+x^{2}+x$$

...

$$Arr CRC4(0x77) = CRC4(0x1D) =$$
 $CRC4(0x48) = CRC4(0x3B) = 0xE$







</math-nerd-session>





Back to BTRS

- Idea: let's fill a directory with files whose names collide under CRC32c, and let's see what happens!
- Demo time



- You can easily fill a bucket (dedicated CRC32c value), and BTRFS will allow you to create only a very limited number of files whose names collide under CRC32c
- You can also easily DoS the bash command line expansion mechanism ;-)!





De Chris Mason

Sujet Re: [btrfs] security hole disclosure

Pour Pascal JUNOD *

Copie à Chris Mason

Hi Pascal,

First, thanks for contacting me and for the time you've spent looking at Btrfs.

Collisions are a known issue with any of the hash based directories. Ext3's tea hash is a little more involved but it is still possible to do the same kinds of collision based DOS attacks. Other 32bit hashes do work better, but even though they are much more CPU intensive, it is still possible to DOS them.

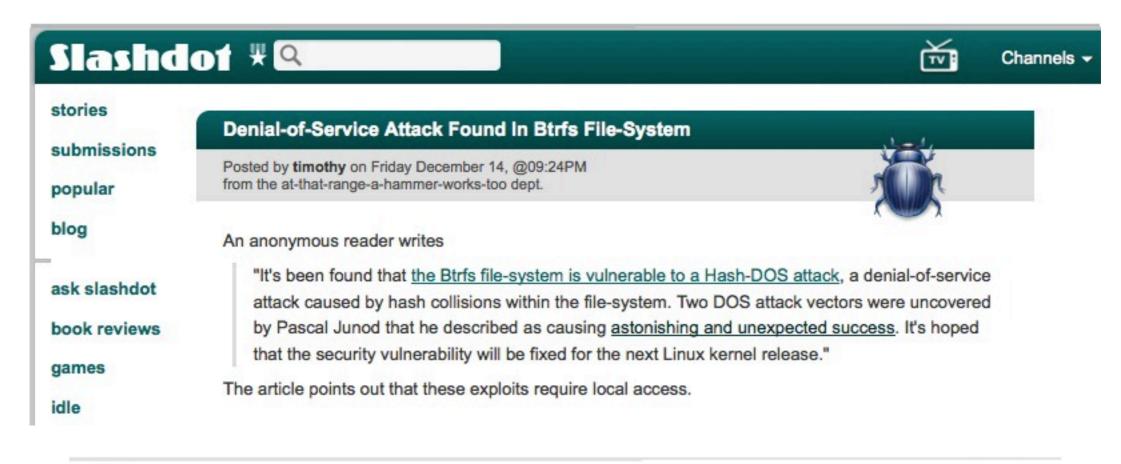
64 bit hashes would be better, but the application interfaces (readdir,seekdir,telldir) on 32 bit machines don't work well with 64 bit directory offsets. Btrfs is able to get around this at least, we return something other than the hash from readdir.

We could make your exploit more complex by salting the hash, but I've always thought the 32 bit hashes were weak enough that salting alone wasn't a huge benefit. If you have any input there I'm very interested in hearing it.

The impact of the DOS is that a malicious user is able to prevent the creation of specific file names. For this to impact other users on the system, it must be done in directories where both the malicious user and the victim user have permission to create files.







On 14 Dec 2012 at 17:14

I strongly suspect this is not a BTRFS issue at all. Perhaps you have tried comparing results with other file systems as well? FWIW, I ran exploits.py on EXT4, XFS and BTRFS formatted partitions and had *no* issues in removing the files. On one trial run with hash=True and ITERATIONS=1000 (i.e.55000 files), I even got more files on BTRFS (54268) than on EXT4.Removing them took less 5 seconds. All this was done on a openSUSE 12.1 machine.

As for the infinite loop theory, rm -f* seems to stall in your case perhaps due to the Bash expansion of the '*'. Try removing the files using xargs ("find . -printo | xargs -o rm")

 $Source: \underline{http://linux.slashdot.org/story/12/12/15/0055217/denial-of-service-attack-found-in-btrfs-file-system}$





De Chris Mason

Sujet Re: [btrfs] looping

Pour Moi <pascal@junod.info> *

Copie à Chris Mason *

Hi Pascal,

You've found a real problem in btrfs though. Changes since I tested the crc overflow handling have made us not deal with our EOVERFLOW error handling completely right, and there are cases where we force the FS readonly when we shouldn't.

So this is a valid bug, I'm just waiting on some review of my fix, which will get backported to a number of kernels.

-chris

On Sat, Dec 15, 2012 at 04:49:02AM -0700, Pascal Junod wrote: Hi Chris,

I have good news for you.

Some people have noticed (in the comments of my blog post) that the CPU is burnt in userland, and not in kernel, a detail that I had unfortunately missed. Some have suggested that actually, btrfs was not looping, but that it was the expansion code of bash. As I generate random filenames, it is likely to have a * or a ? character, and then trigger a complexity attack on the bash command line expansion code, even before going into kernel code.

stracing the rm command has not given anything as output, but I just tried to run my python code on another file-system (ext4), and it loops as well. Argh!

Implementing another way to remove the files (with help of find and xargs feeding rm), I observed an average performance loss of about 10%, which is definitely acceptable. Ironically, the DoS seems to be prevented by the finite number of files mapping to the same key.

Anyway, the filling-bucket "attack" remains still valid, it keeps an open question whether it must be accepted or not. I still tend to think that it is not, as a FS should be robust in every possible situation, including the ones that we don't think about.





Apache





The Apache Case

```
pjunod@fedora:~/software/httpd-2.4.3

File Edit View Search Terminal Help
[pjunod@fedora httpd-2.4.3]$ pwd
/home/pjunod/software/httpd-2.4.3
[pjunod@fedora httpd-2.4.3]$ find . -type f -name "*.c" -exec grep hash {} \; | wc -l

[pjunod@fedora httpd-2.4.3]$ 
[pjunod@fedora httpd-2.4.3]$
```

- Unfortunately, as for nginx, most of the hash tables used in Apache are fed with data coming from configuration files :-(
- Still, one can play a bit!





The Apache Case

```
pjunod@fedora: "/software/httpd-2.4.3/modules/aaa
 File Edit View Search Terminal Help
#define DFLT NONCE LIFE apr time from sec(300)
#define NEXTNONCE DELTA apr time from sec(30)
#define NONCE TIME LEN (((sizeof(apr time t)+2)/3)*4)
#define NONCE HASH LEN (2*APR SHA1 DIGESTSIZE)
#define NONCE LEN (int )(NONCE TIME LEN + NONCE HASH LEN)
#define SECRET LEN
/* client list definitions */
typedef struct hash entry {
                                                /* the key for this entry
    unsigned long
                       key;
                                               /* next entry in the bucket */
    struct hash entry *next;
                                               /* for nonce-count checking */
    unsigned long
                   nonce count;
                      hal[2*APR MD5 DIGESTSIZE+1]; /* for algorithm=MD5-sess
    char
                      last nonce[NONCE LEN+1]; /* for one-time nonce's
    char
  client_entry;
static struct hash table {
    client entry **table;
    unsigned long tbl len;
    unsigned long num entries;
    unsigned long num created;
    unsigned long num removed;
    unsigned long
                   num renewed;
  *client list;
```





MD5-sess

- So it seems that Apache uses a hash table (stored in a shared memory segment) to store data related to the MD5-sess digest authentication mechanism.
- What is MD5-sess?! Only heard about MD5 digest authentication...



MD5-sess

Overview [edit]

Digest access authentication was originally specified by RFC 2069 (An Extension to HTTP: Digest Access Authentication). RFC 2069 (Proposed specifies roughly a traditional digest authentication scheme with security maintained by a server-generated nonce value. The authentication response is formed as follows (where HA1, HA2, A1, A2 are names of string variables):

$$HA1 = MD5(A1) = MD5(username : realm : password)$$

 $HA2 = MD5(A2) = MD5(method : digestURI)$
 $response = MD5(HA1 : nonce : HA2)$

RFC 2069 & was later replaced by RFC 2617 (HTTP Authentication: Basic and Digest Access Authentication). RFC 2617 introduced a number of optional security enhancements to digest authentication; "quality of protection" (qop), nonce counter incremented by client, and a client-generated random nonce. These enhancements are designed to protect against, for example, chosen-plaintext attack cryptanalysis.

If the algorithm directive's value is "MD5" or unspecified, then HA1 is

$$HA1 = MD5(A1) = MD5(username : realm : password)$$

If the algorithm directive's value is "MD5-sess", then HA1 is

$$HA1 = MD5(A1) = MD5(MD5(username : realm : password) : nonce : cnonce)$$

If the gop directive's value is "auth" or is unspecified, then HA2 is

$$HA2 = MD5(A2) = MD5(method : digestURI)$$

If the gop directive's value is "auth-int", then HA2 is

$$HA2 = MD5(A2) = MD5(method : digestURI : MD5(entityBody))$$

If the qop directive's value is "auth" or "auth-int", then compute the response as follows:

$$response = MD5(HA1 : nonce : nonceCount : clientNonce : qop : HA2)$$

If the qop directive is unspecified, then compute the response as follows:

$$response = MD5(HA1 : nonce : HA2)$$

The above shows that when gop is not specified, the simpler RFC 2069 & standard is followed.





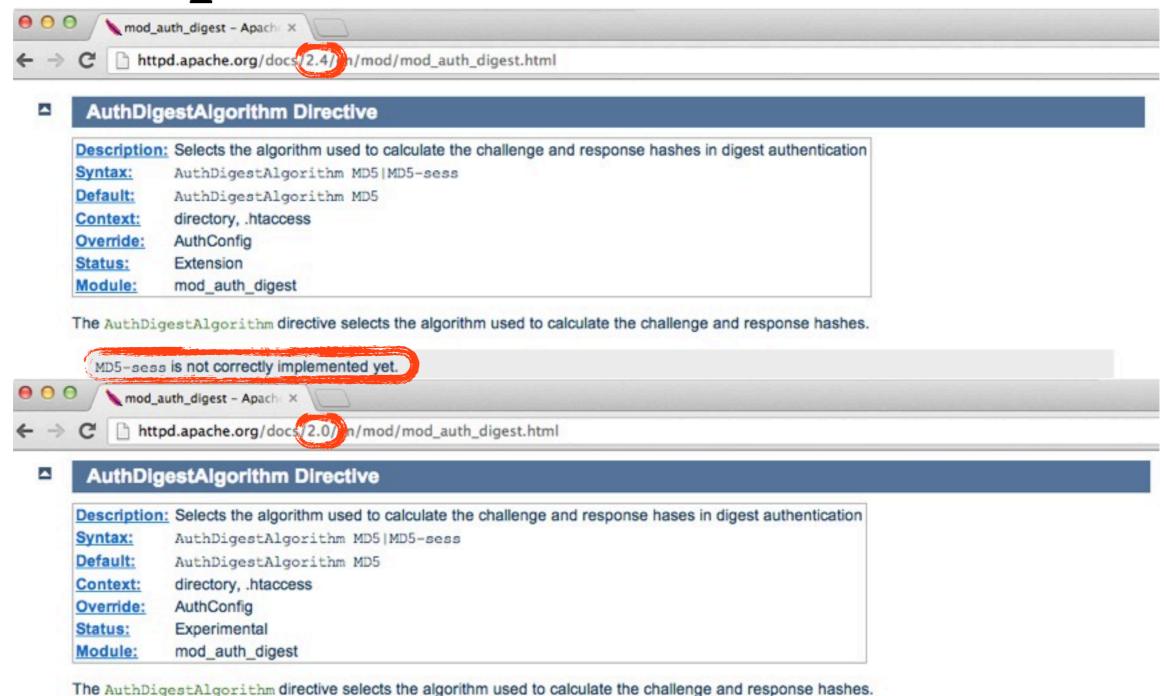
MD5-sess

If the "algorithm" directive's value is "MD5-sess", then A1 is calculated only once - on the first request by the client following receipt of a WWW-Authenticate challenge from the server. It uses the server nonce from that challenge, and the first client nonce value to construct A1 as follows:

This creates a 'session key' for the authentication of subsequent requests and responses which is different for each "authentication session", thus limiting the amount of material hashed with any one key. (Note: see further discussion of the authentication session in section 3.3.) Because the server need only use the hash of the user credentials in order to create the A1 value, this construction could be used in conjunction with a third party authentication service so that the web server would not need the actual password value. The specification of such a protocol is beyond the scope of this specification.







MD5-sess is not correctly implemented yet.



```
[pjunod@fedora aaa]$ pwd
/home/pjunod/local/httpd/htdocs/aaa
[pjunod@fedora aaa]$ cat .htaccess
AuthType Digest
AuthDigestAlgorithm MD5-sess
AuthName "test"
AuthDigestProvider file
AuthUserFile /home/pjunod/local/httpd/htdocs/
aaa/.htpasswd
AuthDigestNonceLifetime 0
Require user pjunod
```





```
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.
GET /aaa HTTP/1.0
HTTP/1.1 500 Internal Server Error
Date: Thu, 21 Mar 2013 14:25:58 GMT
Server: Apache/2.4.3 (Unix)
Content-Length: 528
Connection: close
Content-Type: text/html; charset=iso-8859-1
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>500 Internal Server Error</title>
</head><body>
<h1>Internal Server Error</h1>
The server encountered an internal error or
misconfiguration and was unable to complete
your request.
Please contact the server administrator at
 you@example.com to inform them of the time this error occurred,
 and the actions you performed just before this error.
More information about this error may be available
in the server error log.
</body></html>
Connection closed by foreign host.
```





[pjunod@fedora aaa]\$ telnet localhost 80

```
[pjunod@fedora aaa]$ tail -n1 ~/local/httpd/logs/error_log
[Thu Mar 21 15:25:59.858082 2013] [core:alert] [pid 23526:tid
139779710646016] [client 127.0.0.1:59616] /home/pjunod/local/httpd/
htdocs/aaa/.htaccess: AuthDigestAlgorithm: ERROR: algorithm `MD5-
sess' is not fully implemented
```





```
Opaque and hash-table management
A string of
data
                 * Generate a new client entry, add it to the list, and return the
                 * entry. Returns NULL if failed.
specified by
                static client entry *gen client(const request rec *r)
the Server
                   unsigned long op;
that
                    client entry new entry = { 0, NULL, 0, "", "" }, *entry;
contains a
                   if (!opaque cntr) {
                        return NULL;
reference
for the
                    apr global mutex lock(opaque lock);
                    op = (*opaque cntr)++;
Security
                    apr global mutex unlock(opaque lock);
context
                    if (!(entry = add client(op, &new entry, r->server))) {
                        ap log rerror(APLOG MARK, APLOG ERR, 0, r, APLOGNO(01769)
that is
                                     "failed to allocate client entry - ignoring client");
                        return NULL;
being
established.
                    return entry;
```





- In summary:
 - Craft an HTTP query with
 - an Authorization field of type Digest MD5
 - no opaque directive
 - a AuthDigestNonceLifetime 0 directive server-side
- and let's see what happens!





```
[Fri Mar 22 19:16:25.832503 2013] [core:notice] [pid 23524:tid 139779852523264] AH00052: child pid 29275 exit signal Floating point exception (8) [Fri Mar 22 19:16:25.832554 2013] [core:notice] [pid 23524:tid 139779852523264] AH00052: child pid 29276 exit signal Floating point exception (8) [Fri Mar 22 19:16:25.832561 2013] [core:notice] [pid 23524:tid 139779852523264] AH00052: child pid 29277 exit signal Floating point exception (8) [Fri Mar 22 19:16:25.832568 2013] [core:notice] [pid 23524:tid 139779852523264] AH00052: child pid 29278 exit signal Floating point exception (8)
```





- In other words, any user can sabotage an Apache installation in a shared web environment
- Possibly, this can be transformed in a remote-only attack if one is able to upload an .htaccess file in a way or in another.



De Mark J Cox

Sujet Re: security hole disclosure

Pour Pascal JUNOD

Copie à Apache Software Foundation HTTP Server Project

[resent, was missing cc]

I apologise for the delay in responding to you; we give priority to critical and important issues and so lower severity issues and bugs tend to keep getting pushed to the bottom of the pile. In general we would not treat an issue like this as a security threat; it requires a local malicious user; and if your local attacker can craft a .htaccess file then there are likely many other ways they could do so which could cause an Apache DoS. the next step would be to send this as a bug report (or to httpd-dev) for discussion.

Regards, Mark

On Wed, Dec 5, 2012 at 1:37 PM, Pascal Junod pascal.junod@heig-vd.ch wrote: Dear Apache security team,

He says
it's a
banana
attack:-(











- Still, there is a morale attached to this story:
 - When you have unused code in your project, don't compile it!







Contact

http://crypto.junod.info
pascal@junod.info
@cryptopathe



