MAGAZINE BSD

FOR NOVICE AND ADVANCED USERS

NodeJS and FreeBSD

HOW TO BUILD NODEJS FROM SOURCE CODE ON FREEBSD

ZFS POOL CONFIGURATION

INFORMATION SECURITY

BASECAMP – PROJECT MANAGEMENT FOR THE SANE

VOL.9 NO.04
ISSUE 69
1898-9144

FREENAS MINI
STORAGE APPLIANCE

IT *SAVES* YOUR LIFE.

HOW IMPORTANT IS YOUR DATA?

Years of family photos. Your entire music and movie collection. Office documents you've put hours of work into. Backups for every computer you own. We ask again, *how important is your data?*

NOW IMAGINE LOSING IT ALL

Losing one bit - that's all it takes. One single bit, and your file is gone.

The worst part? **You won't know until you absolutely need that file again.**

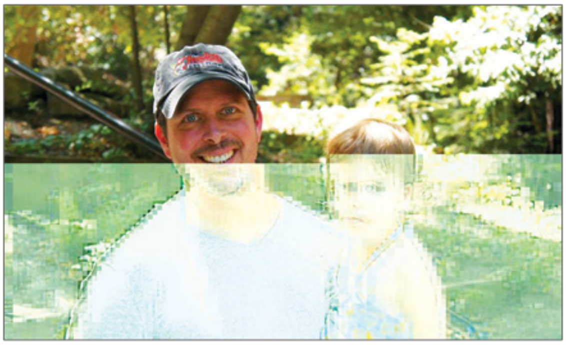

Example of one-bit corruption

THE SOLUTION

The FreeNAS Mini has emerged as the clear choice to save your digital life. **No other NAS in its class offers ECC (error correcting code) memory and ZFS bitrot protection to ensure data always reaches disk without corruption and *never degrades over time*.**

No other NAS combines the inherent data integrity and security of the ZFS filesystem with fast on-disk encryption. No other NAS provides comparable power and flexibility. The FreeNAS Mini is, hands-down, the best home and small office storage appliance you can buy on the market. **When it comes to saving your important data, there simply is no other solution.**

The Mini boasts these state-of-the-art features:

- 8-core 2.4GHz Intel® Atom™ processor
- Up to 16TB of storage capacity
- 16GB of ECC memory (with the option to upgrade to 32GB)
- 2 x 1 Gigabit network controllers
- Remote management port (IPMI)
- Tool-less design; hot swappable drive trays
- FreeNAS installed and configured

iXsystems®

http://www.iXsystems.com/mini

intel® inside™ ATOM™

FREENAS CERTIFIED
STORAGE

With over six million downloads, FreeNAS is undisputedly *the* most popular storage operating system in the world.

Sure, you could build your own FreeNAS system: research every hardware option, order all the parts, wait for everything to ship and arrive, vent at customer service because it *hasn't*, and finally build it yourself while hoping everything fits - only to install the software and discover that the system you spent *days* agonizing over **isn't even compatible**. Or...

MAKE IT EASY ON YOURSELF

As the sponsors and lead developers of the FreeNAS project, iXsystems has combined over 20 years of hardware experience with our FreeNAS expertise to bring you FreeNAS Certified Storage. **We make it easy to enjoy all the benefits of FreeNAS without the headache of building, setting up, configuring, and supporting it yourself.** As one of the leaders in the storage industry, you know that you're getting the best combination of hardware designed for optimal performance with FreeNAS.

Every FreeNAS server we ship is...

» Custom built and optimized for your use case
» Installed, configured, tested, and guaranteed to work out of the box
» Supported by the Silicon Valley team that designed and built it
» Backed by a 3 years parts and labor limited warranty

As one of the leaders in the storage industry, you know that you're getting the best combination of hardware designed for optimal performance with FreeNAS. **Contact us today for a FREE Risk Elimination Consultation with one of our FreeNAS experts.** Remember, every purchase directly supports the FreeNAS project so we can continue adding features and improvements to the software for years to come. **And really - why would you buy a FreeNAS server from *anyone* else?**

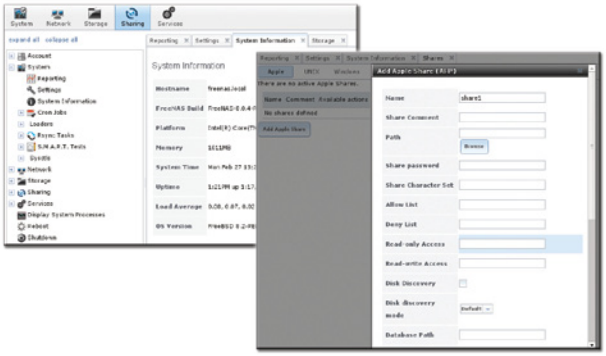

System Information

Hostname	freenas.local
FreeNAS Build	FreeNAS-8.0.4
Platform	Intel(R) Core(T
Memory	1011MB
System Time	Mon Feb 27 13:
Uptime	1:21PM up 1:17
Load Average	0.08, 0.07, 0.0
OS Version	FreeBSD 8.2-R

There are no active Apple Shares.

Name Comment Available actions
No shares defined

Name	share1
Share Comment	
Path	
Share password	
Share Character Set	
Allow List	
Deny List	
Read-only Access	
Read-write Access	
Disk Discovery	□
Disk discovery mode	Default
Database Path	

FreeNAS 1U

- Intel® Xeon® Processor E3-1200v2 Family
- Up to 16TB of storage capacity
- 16GB ECC memory (upgradable to 32GB)
- 2 x 10/100/1000 Gigabit Ethernet controllers
- Redundant power supply

FreeNAS 2U

- 2x Intel® Xeon® Processors E5-2600v2 Family
- Up to 48TB of storage capacity
- 32GB ECC memory (upgradable to 128GB)
- 4 x 1GbE Network interface (Onboard) - (Upgradable to 2 x 10 Gigabit Interface)
- Redundant Power Supply

intel inside™ XEON®

http://www.iXsystems.com/storage/freenas-certified-storage/

Dear Readers,

The BSD magazine team is pleased to announce the launch of the next issue of BSD Magazine.

A lot of tutorials and practice rich articles are included in this issue to help you develop your skills and knowledge. Our ultimate goal is to provide our readers with exactly the knowledge and skills they need in their IT careers. Hence, we will be very glad to receive your suggestions for workshops, tutorials, what you need most, etc…

Let's take a look at what you will learn in this issue.

Our experts will teach you how to build nodejs from source code on FreeBSD. In addition, you will discover ZFS Pool Configuration and how to create a RAIDZ2 of ten drives.

In addition, you will will develop a basic understanding of the project management tool, Basecamp, as well as learn how to get up to speed quickly with Basecamp so that you can start realizing its benefits..

We wish to say "Thank You" and express our gratitude to our experts who contributed to this issue and to our coming issues. We invite other experts for collaboration for the next issue, due out in 4 weeks.

Stay tuned, we have two special issues that will be published soon.

Enjoy reading,
Ewa & BSD Team

MAGAZINE BSD

Editor in Chief:
Ewa Dudzic
ewa.dudzic@software.com.pl

Contributing:
Michael Shirk, Andrey Vedikhin, Petr Topiarz, Solène Rapenne, Anton Borisov, Jeroen van Nieuwenhuizen, José B. Alós, Luke Marsden, Salih Khan, Arkadiusz Majewski, BEng, Toki Winter, Wesley Mouedine Assaby, Rob Somerville

Top Betatesters & Proofreaders:
Annie Zhang, Denise Ebery, Eric Geissinger, Luca Ferrari, Imad Soltani, Olaoluwa Omokanwaye, Radjis Mahangoe, Mani Kanth, Ben Milman, Mark VonFange

Special Thanks:
Annie Zhang
Denise Ebery

Art Director:
Ireneusz Pogroszewski

DTP:
Ireneusz Pogroszewski
ireneusz.pogroszewski@software.com.pl

Senior Consultant/Publisher:
Paweł Marciniak
pawel@software.com.pl

CEO:
Ewa Dudzic
ewa.dudzic@software.com.pl

Publisher:
Hakin9 Media SK
02-676 Warsaw, Poland
Postepu 17D
Poland
worldwide publishing
editors@bsdmag.org
www.bsdmag.org

Hakin9 Media SK is looking for partners from all over the world. If you are interested in cooperation with us, please contact us via e-mail: editors@bsdmag.org.

FreeNAS
in an Enterprise Environment

By the time you're reading this, FreeNAS has been downloaded more than 5.5 million times. For home users, it's become an indispensable part of their daily lives, akin to the DVR. Meanwhile, all over the world, thousands of businesses universities, and government departments use FreeNAS to build effective storage solutions in myriad applications.

What you will learn...

- How TrueNAS builds off the strong points of the FreeBSD and FreeNAS operating systems
- How TrueNAS meets modern storage challenges for enterp

The FreeNAS operating systems is fre
the public and offers thorough doc
active community, and a feature-ric
the storage environment. Based on Free
can share over a host of protocols (SMB
FTP, iSCSI, etc) and features an intuitiv
the ZFS file system, a plug-in system
much more.

Despite the massive popularity o
aren't aware of its big brother duti
data in some of the most demand
environments: the proven, enterp
professionally-supported line of

But what makes TrueNAS diffe
Well, I'm glad you asked...

Commercial Grade Supp

When a mission critical stor
organization's whole operat
halt. Whole community-bas
free), it can't always get an
and running in a timely m
responsiveness and expe
dedicated support team
provide that safety.

Created by the sam
developed FreeNAS.

FreeBSD and NodeJS

Project management

Expert says...

Security Corner

Other Technologies

Column

SONY
riis
AnDevCon
Exhibit H
intel inside

NodeJS and FreeBSD – Part 1

DAVID CARLIER

Nodejs is well known to allow building server applications in full javascript. In this article, we'll see how to build nodejs from source code on FreeBSD. You will need autoconf tools, GNU make, Python, linprocfs enabled and libexecinfo installed. GCC/G++ compiler suite (C++11 compliant, ideally 4.8 series or above) or possibly clang can be used to compile the whole source.

To start, we need the nodejs source code from this url *http://www.nodejs.org/dist/latest* where we can find this archive (during the article writing, the last version known is 0.12.2), `node-v<version>.tar.gz`.

Be prepared to be patient, you have enough time for a cup of coffee, the compilation time needed can be quite long...

Once downloaded and extracted, the famous command trio needs to be typed:

- ./configure --dest-os=freebsd
- gmake
- gmake install

It's pretty straightforward on first glance. On FreeBSD, when v8 is compiled we get some compilation errors:

```
clang++ '-DV8_TARGET_ARCH_X64' '-DENABLE_DISASSEMBLER'
    '-DENABLE_HANDLE_ZAPPING' -I../deps/v8  -pthread
    -Wall -Wextra -Wno-unused-parameter -m64 -fno-strict-
    aliasing -I/usr/local/include -O3 -ffunction-sections
    -fdata-sections -fno-omit-frame-pointer -fdata-sections
    -ffunction-sections -O3 -fno-rtti -fno-exceptions -MMD
    -MF /root/node-v0.12.2/out/Release/.deps//root/node-
    v0.12.2/out/Release/obj.target/v8_libbase/deps/v8/src/
    base/platform/platform-freebsd.o.d.raw  -c -o /root/
    node-v0.12.2/out/Release/obj.target/v8_libbase/deps/v8/
    src/base/platform/platform-freebsd.o ../deps/v8/src/
    base/platform/platform-freebsd.cc
../deps/v8/src/base/platform/platform-freebsd.cc:159:11:
    error: member reference base type 'int' is not a
    structure or union
     result.push_back(SharedLibraryAddress(start_of_path,
    start, end));
     ~~~~~~^~~~~~~~~~
../deps/v8/src/base/platform/platform-freebsd.cc:191:53:
    error: use of undeclared identifier 'MAP_NORESERVE'
                            MAP_PRIVATE | MAP_ANON | MAP_
    NORESERVE,
                                                     ^
../deps/v8/src/base/platform/platform-freebsd.cc:263:48:
    error: use of undeclared identifier 'MAP_NORESERVE'
                       MAP_PRIVATE | MAP_ANON | MAP_
    NORESERVE,
                                                ^
../deps/v8/src/base/platform/platform-freebsd.cc:291:40:
    error: use of undeclared identifier 'MAP_NORESERVE'
```

```
              MAP_PRIVATE | MAP_ANON | MAP_NORESERVE |
   MAP_FIXED,
                                            ^
```

4 errors generated.

Ok, so a result variable ought to be a std::vector but it’s considered wrongly as an int and furthermore a wrong mmap flag is used. Let’s fix it!

```
std::vector<SharedLibraryAddress> result;
  static const int MAP_LENGTH = 1024;
  int fd = open(“/proc/self/maps”, O_RDONLY);
  if (fd < 0) return result;
  while (true) {
    char addr_buffer[11];
    addr_buffer[0] = ‘0’;
    addr_buffer[1] = ‘x’;
    addr_buffer[10] = 0;
    int result = read(fd, addr_buffer + 2, 8);
    if (result < 8) break;
    unsigned start = StringToLong(addr_buffer);
    result = read(fd, addr_buffer + 2, 1);
    if (result < 1) break;
    if (addr_buffer[2] != ‘-’) break;
    result = read(fd, addr_buffer + 2, 8);
    if (result < 8) break;
    unsigned end = StringToLong(addr_buffer);
    char buffer[MAP_LENGTH];
    int bytes_read = -1;
   do {
     bytes_read++;
     if (bytes_read >= MAP_LENGTH - 1)
       break;
     result = read(fd, buffer + bytes_read, 1);
```

Apparently, there are two different variables with the same name. Let’s rename the second, the int type, to res for example so the vector result variable can legitimately call push_back method. That fixes the first error.

```
std::vector<SharedLibraryAddress> result;
  static const int MAP_LENGTH = 1024;
  int fd = open(“/proc/self/maps”, O_RDONLY);
  if (fd < 0) return result;
  while (true) {
    char addr_buffer[11];
    addr_buffer[0] = ‘0’;
    addr_buffer[1] = ‘x’;
    addr_buffer[10] = 0;
    int res= read(fd, addr_buffer + 2, 8);
    if (res < 8) break;
    unsigned start = StringToLong(addr_buffer);
    res = read(fd, addr_buffer + 2, 1);
    if (res < 1) break;
    if (addr_buffer[2] != ‘-’) break;
    res = read(fd, addr_buffer + 2, 8);
    if (res < 8) break;
    unsigned end = StringToLong(addr_buffer);
    char buffer[MAP_LENGTH];
    int bytes_read = -1;
    do {
      bytes_read++;
      if (bytes_read >= MAP_LENGTH - 1)
        break;
      res = read(fd, buffer + bytes_read, 1);
```

Let’s have a look at the mmap problem.

`MAP_NORESERVE` is a specific flag which guarantees no swap space will be used for the mapping. However, it is a flag usable on Linux and Solaris /SunOS.

```
mmap(OS::GetRandomMmapAddr(),
                      size,
                      PROT_NONE,
                      MAP_PRIVATE | MAP_ANON | MAP_
   NORESERVE,
                      kMmapFd,
                      kMmapFdOffset);

=>

mmap(OS::GetRandomMmapAddr(),
                      size,
                      PROT_NONE,
                      MAP_PRIVATE | MAP_ANON,
                      kMmapFd,
                      kMmapFdOffset);


void* reservation = mmap(OS::GetRandomMmapAddr(),
                           request_size,
                           PROT_NONE,
                           MAP_PRIVATE | MAP_ANON | MAP_
   NORESERVE,
                           kMmapFd,
                           kMmapFdOffset);

=>

void* reservation = mmap(OS::GetRandomMmapAddr(),
                           request_size,
                           PROT_NONE,
                           MAP_PRIVATE | MAP_ANON,
                           kMmapFd,
                           kMmapFdOffset);
```

Once modified in every mmap call, we can now retry compiling. However, we get another compilation error. This time, it casts a pthread_self returns call to an int.

```
deps/v8/src/base/platform/platform-posix.cc:331:10: error:
   static_cast from ‘pthread_t’ (aka ‘pthread *’) to ‘int’
   is not allowed
  return static_cast<int>(pthread_self());
```

The problem is, on FreeBSD, a pthread_t type is not an integral type at all but an opaque struct.

Instead, we might replace this line by:

```
return static_cast<int>(reinterpret_cast<inptr_t>(pthread_
   self()));
```

Now we are finally able to compile. After a couple of minutes, it is finished but we have still one source to update: lib/dns.js. Add these two lines after line 127:

```
if (process.platform === ‘freebsd’ && family !== 6)
        hints &= ~exports.V4MAPPED;
```

Because FreeBSD does not support this flag, it ought to be cleared. This is all for compilation and it is ready to be used. Next time, we’ll have an overlook in the application’s building part and ought to see the potential of this library.

ABOUT THE AUTHOR

David Carlier has been working as a software developer since 2001. He used FreeBSD for more than 10 years and starting from this year, he became involved with the HardenedBSD project and performed serious developments on FreeBSD. He worked for a mobile product company that provides C++ APIs for two years in Ireland. From this, he became completely inspired to develop on FreeBSD.

ISO

mobile · interactive · design

- ☑ Mobile Apps
- ☑ Website Design
- ☑ Specialty Programming
- ☑ 3DSimulations
- ☑ Unity 3D
- ☑ SmartFoxServer
- ☑ Games
- ☑ Web & Database Dev
- ☑ Super friendly :)

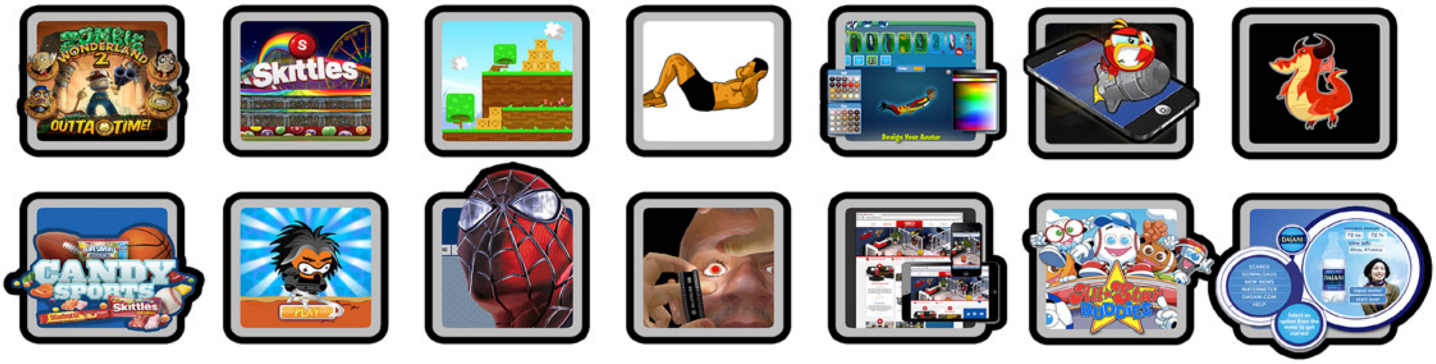

reach out & let's talk: troy@isointeractive.com

www.isointeractive.com

Basecamp – Project Management for the Sane

TROY HIPOLITO

Download the latest ISO Interactive white paper. There you will find a company description, capabilities, visuals, development process, case insights, and technology definitions.

- ISO White Paper: *www.isointeractive.com/pdf*
- ISO Video: *www.isointeractive.com/#showreel*
- ISO Website: *www.isointeractive.com*

ISO Interactive are award winning consultants that build engaging mobile and web experiences. Known for small to large opportunities using Unity, Flash, HTML5 and traditional web programming, they have built very cool virtual worlds, 3D simulations, mobile apps, social games and web designs.

Overview

In this tutorial, we will dive into a basic understanding of Basecamp (a project management tool we use), as well as learn how to get up to speed quickly so that you can start realizing the benefits of the program, among which are centralizing communications, reducing the frequency of meetings, facilitating team coordination on projects, and providing transparency on timelines.

We do have more detailed information concerning the project management role and methods that work best for your orginization in my previous article located at: *http://sdjournal.org/download/2011-pentest-extra-issues/*. Feel free to check it out as there is good information on project management organization and methods.

Speaking of... Project management is one area we have a lot of experience in. We believe project management is a major factor in determining success of the project. This is especially true for complex and technical endeavors.

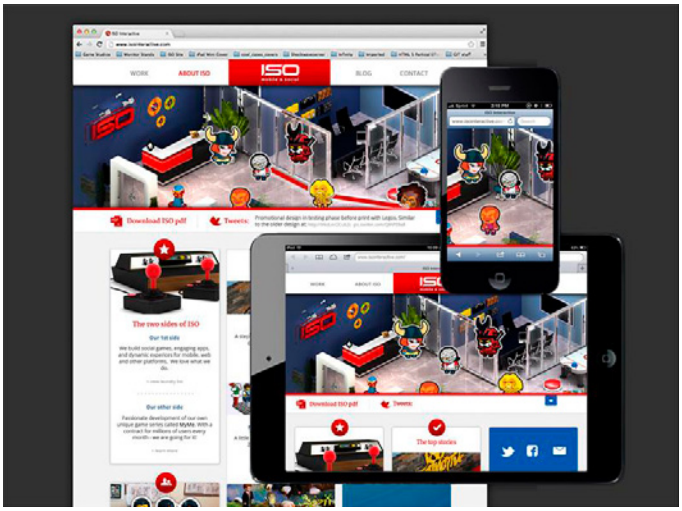

Now I am not taking away from the great designers and developers, but having these is more of a norm. Great designers and developers need unification and sometimes direction to keep goals, budgets and timelines reasonable.

Our groups have worked in corporate as well as the agency scenarios. To be honest, we favor an agency style as it has more of a startup feel and allows us to get our hands dirty. This allows some control to drive tasks and better target success.

Corporate project management is our view more in reporting to a number of bosses than actual management. It's different due to the structure and size of the client/partner.

The good people at 37signals have revamped their popular project management software Basecamp. Previously we produced a popular project management article for the Software Developer's Journal that touches on the old version of the software. More specifically, it is the cover article for the Flash & Flex magazine in 2011.

So we have actually touched on some of that information but now we will concentrate on an in-depth tutorial of the new version of Basecamp.

This tutorial is divided into several sections, starting with the basic Why Basecamp?, followed by a description of the various features and capabilities of Basecamp. The third section will cover usage instructions and guidelines, from identifying project scope to replying to Basecamp Messages. The final section covers the conclusion.

Why Basecamp?

You may be wondering, *why do we use Basecamp versus another tool?*

Well we actually do use other tools depending on the client/partner/requirements. There are many great online tools out there, for example Jira, MS Project, Asana and RallyDev. Some of these are more feature rich with true Agile processes while others have very specific set of functions.

At ISO the main focus is to produce a high quality product with the least amount of drama. That may not sound completely intuitive, but if you think about it, everything is about making things flow and reducing drama. Controlling costs is actually a byproduct.

The best designers and developers are sometimes a pain in the butt (not all, but most). You know what I'm talking about: acting like they just hit puberty, not making their deadlines (that they committed to), getting their feelings hurt easily, whining, crying and all that nonsense. And they have to be managed without them pooping their pants and walking out of the job because they aren't doing what they said or aren't getting their way. My goodness, it is pain to manage but absolutely needed.

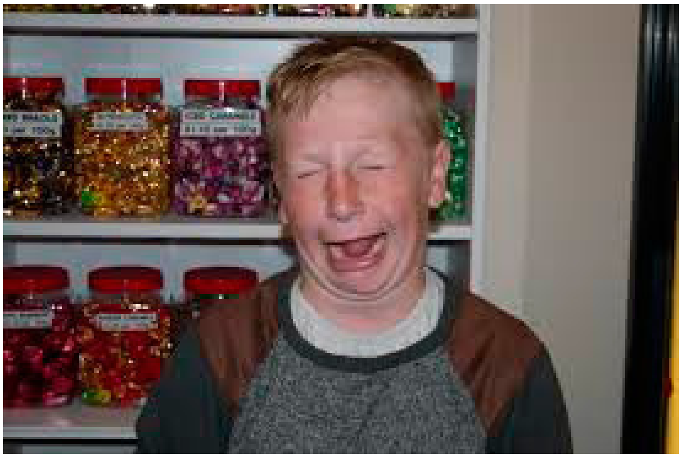

While generally we "try" to adopt more agile processes, we are bound by the rapid changing needs of the business, which can grow in volume at a rapid pace. Our focus is directed by numerous initiatives that result in a compound of projects with a pairing of unique groups. Planning projects around Agile-style "sprints" (i.e., a guaranteed amount of time) is not always possible and more often not probable.

Basecamp is often more suitable for many of our needs because it is task-oriented and date-driven. Another great benefit of Basecamp is it's an entirely online secure desktop tool. Basecamp also offers a mobile app. Additional highlights the program offers:

- Centralizing communication for emails based on the project, conversation thread and assigned tasks.
- Uploading and tagging files associated with a particular project. Typically, these are items like word documents, spreadsheets, images, PSDs and PDFs.
- Setting up and tracking schedules for development, meetings, and handoffs.

The key to success when using Basecamp is for everyone to actually use it for the tasks at hand. Otherwise, there will not be record of any tasks being worked on. This can easily degenerate into halting progression to the next step of the project, delays in securing approvals and handing off to other departments, and failure to meet deadlines. In short, not properly communicating within Basecamp and your project tasks can jeopardize launch dates.

So think of Basecamp as a handy organization tool that allows the your team to be more efficient and enhance productivity.

On to our review of Basecamp!

The following are the six (6) main sections found under the Projects Menu in Basecamp:

- Projects
- Calendar
- Everything
- Progress
- Everyone
- Me

Projects Menu

When you log in to Basecamp, you are directed to the main page, where you are able to see all of the projects available. From here you can *select the project you want*, *change the view of how you want to see projects*, *"star" projects that pertain to you*, and even *create new projects via a template or from scratch*. On this page there is also a little search box that allows you to find things quickly.

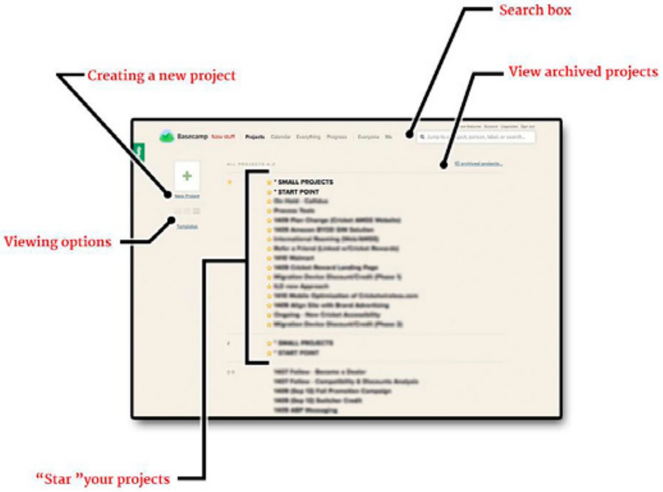

There are a number of projects in the queue at any given time. To find a particular project simply, scroll up or down until find what you need. You'll also have the option to change the view from "graphical" to "hybrid" to "textual" by using the icons on the left of the screen (below the New Project link). If you like to read through the list quickly you may want to use the "textual" view.

If you want to group the projects which are specifically assigned to you, simply click on the "star" for those projects and they will all be moved up together to the top.

Additionally, if you want to look for a project that you know is finished but can't find the name, click the archived projects link on the top right to see a listing of those projects.

Individual Project "Project Name"

Each project has a number of components. Clicking on a project, you will notice menu links/sub sections for Project Landing Page, Discussions, To-dos, Files, Text Documents and Events.

Project Landing Page

The title of the project is the link to the project landing page. These pages are useful for viewing recent activity on the other subsections. From top to bottom it has the Latest project updates, Discussions, To-do lists, a visual of the Files uploaded and newest Text Documents.

All content from these subsections displayed on the project landing page link directly to those details.

Next, let's take a look at the definition of most of these titles.

Discussions

Discussions serves as a type of "centralized" email inbox. Typically Discussions are not directly tied to to-dos.

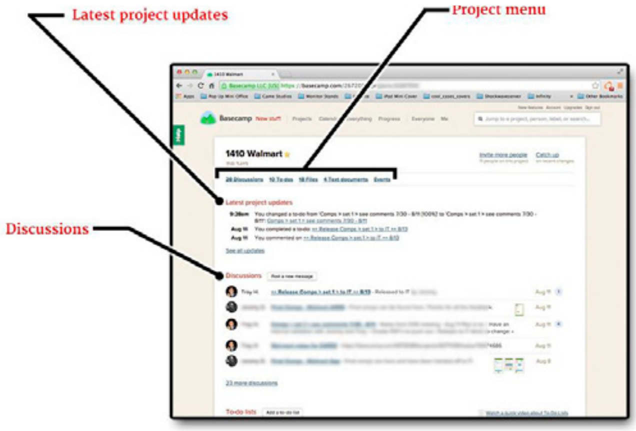

These are used for emails which may deal with internal approvals and general notes that may or may not pertain to current tasks. Each discussion can have its own thread. For example, "discussions" can be used to post project notes.

Starting a Discussion:

To start a discussion, log in to your Basecamp project, click on the post a new message button. There you have the *Subject line* and the *message area*. You can *format* with the tools available and if needed *upload a file*.

Make sure you DO NOT email everyone. No one likes spam. Please only click on the individuals that need to know. There is more on this and other etiquette items in the How We Use Basecamp section.

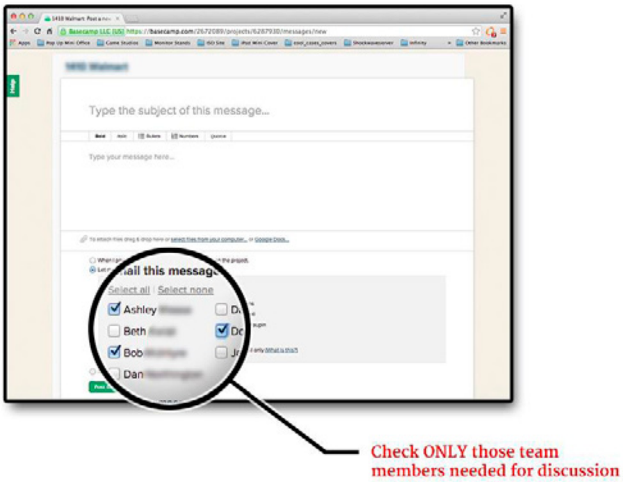
Type the subject of this message...
Type your message here...
Email this message
Select all | Select none
☑ Ashley
□ Beth
☑ Bob
□ Dan
Check ONLY those team members needed for discussion

Once they receive a Basecamp message, individuals can simply email back from their native email client or click the "view on Basecamp" link within the email to reply. Viewing from Basecamp will allow the entire conversation thread to be reviewed. Lastly, you can attach files as needed in discussions.

To-do lists

To-dos are a vital part of Basecamp. It is what generates the Calendar and assigns tasks to individuals and denotes important events. It is basically an adjustable task list with due dates.

Within the following image there is an *Add a to-do list* button, *title* of a current To-do list and individual *to-dos/tasks*. And on the right side there are the view options (show assigned to, show when is due, show completed, and individual to-do lists).

Basecamp allows for numerous lists. Typically, depending on the type and size of the project, you may want to break it up. At this time, however, our projects are fairly small, so we would we prefer a more linear approach. It is simpler for our current needs.

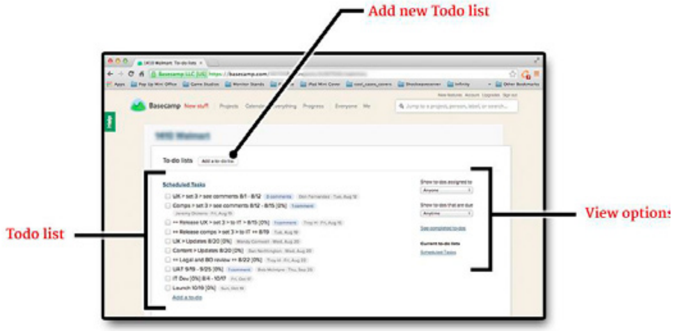

The previous image displays individual to-dos that can have a few pieces of additional information.

Normally it has a description, due date and the person the task is assigned to. If there are any comments relating to this to-do task, you will see a note following the to-do task description.

Once the to-do task is completed, the assigned person or Project Manager can check it off (with the check button).

In order to view a comment, just click on it. Comments are very good if the short description does not have enough detail.

However, for our projects we add some additional info. In each of our to-do items we have the subject (e.g. Comps), short description, percentage complete, date or date range, separate due date and person the task is assigned to.

In the example below, the original description was reduced to UX > set 1. In this case, we added comments for further clarification.

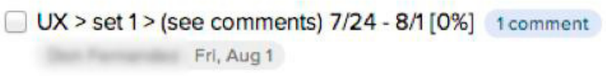

□ UX > set 1 > (see comments) 7/24 - 8/1 [0%] 1 comment
[redacted] Fri, Aug 1

Selecting on a to-do task (that has comments) will display all the comments in a thread fashion. This is very similar to how the message conversations are done. The viewer will also be able to select to whom to email the message. The purpose is to have the person assigned to be responsive regarding the progress of their task and to centralize related conversations. Only conversations that pertain to this to-do task should be added here.

There is also an option to email people outside of Basecamp who are not part of the project. However, that is not recommended for our production flow.

Once individuals receive a message, they can simply reply via their email client or from Basecamp. It is usually better to email back via Basecamp if you feel that you want to read part or the entire thread.

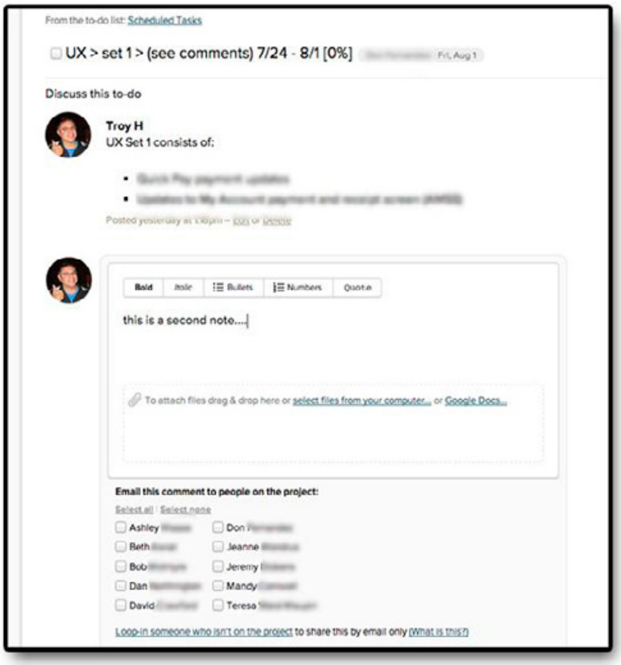

From the to-do list: Scheduled Tasks

□ UX > set 1 > (see comments) 7/24 - 8/1 [0%] [redacted] Fri, Aug 1

Discuss this to-do

Troy H
UX Set 1 consists of:

- [redacted]
- [redacted]

[redacted]

this is a second note....

To attach files drag & drop here or select files from your computer... or Google Docs...

Email this comment to people on the project:
Select all | Select none

□ Ashley [redacted]	□ Don [redacted]
□ Beth [redacted]	□ Jeanne [redacted]
□ Bob [redacted]	□ Jeremy [redacted]
□ Dan [redacted]	□ Mandy [redacted]
□ David [redacted]	□ Teresa [redacted]

Loop-in someone who isn't on the project to share this by email only (What is this?)

Going back to the additional info we added, let's talk about the percentage item (e.g. [30%]). This info can be added manually in Basecamp by the person performing the task, as well as the project manager. This is done by moving your cursor over the to-do task and then selecting the "edit" option.

This allows everyone to quickly see how much of the task is considered completed. Formatting it this way also has some major advantages. That piece of information, along with the due date, is pulled in a visual Gantt chart/timeline called TeamGantt..

Gantt Chart

This is separate online software that is very useful for visual teams. And the project manager can invite the same people on the project from Basecamp to TeamGantt.

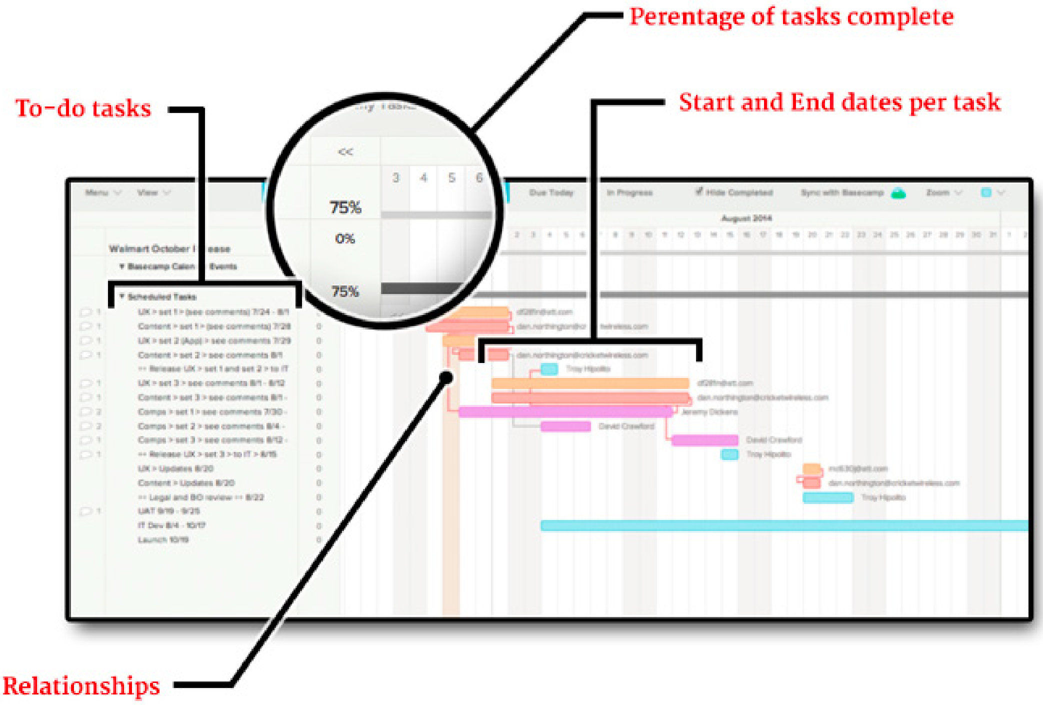
Perentage of tasks complete
To-do tasks
Start and End dates per task
Relationships
75%
0%
75%
Menu
View
Due Today
In Progress
Hide Completed
Sync with Basecamp
Zoom
August 2014
Walmart October Release
Basecamp Calendar Events
Scheduled Tasks
UX > set 1 > (see comments) 7/24 - 8/1
Content > set 1 > (see comments) 7/28
UX > set 2 (App) > see comments 7/29
Content > set 2 > see comments 8/1
>> Release UX > set 1 and set 2 > to IT
UX > set 3 > see comments 8/1 - 8/12
Content > set 3 > see comments 8/1 -
Comps > set 1 > see comments 7/30 -
Comps > set 2 > see comments 8/4 -
Comps > set 3 > see comments 8/12 -
>> Release UX > set 3 > to IT > 8/15
UX > Updates 8/20
Content > Updates 8/20
>> Legal and BO review >> 8/22
UAT 9/19 - 9/25
IT Dev 8/4 - 10/17
Launch 10/19
Troy Hipolito
Jeremy Dickens
David Crawford

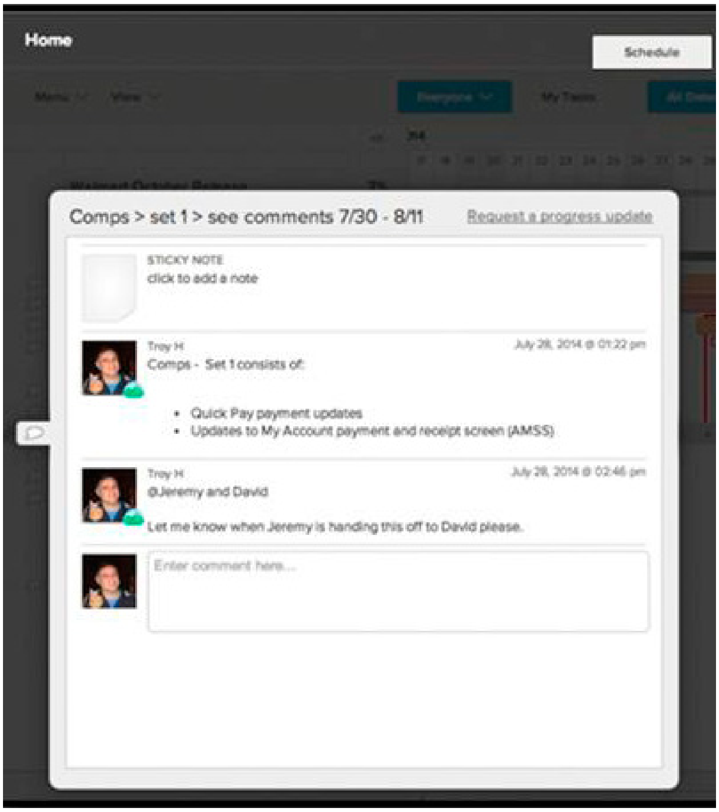

Comps > set 1 > see comments 7/30 - 8/11

Request a progress update

STICKY NOTE
click to add a note

Troy H
July 28, 2014 @ 01:22 pm
Comps - Set 1 consists of:

- Quick Pay payment updates
- Updates to My Account payment and receipt screen (AMSS)

Troy H
July 28, 2014 @ 02:46 pm
@Jeremy and David

Let me know when Jeremy is handing this off to David please.

Enter comment here...

TeamGantt uses the percentage info provided to show completion of the task. The beginning and end dates in the to-do task description are just a reference so we can visually adjust the timeline. But adding the time range in the textual format makes it is easy to read and that is the important thing.

TeamGantt also has some neat printing features, associating tasks with each other and even color coding groups of tasks.

As such, if designing comps are dependent on wireframes, you can link them together visually. In this case, I have made all the comp related items a fuchsia, or "hot pink" color. So all the comp driven tasks are one color, UX-related another color, and so on.

A couple of other neat features allow the Project Managers or individuals with edit permissions to send notes to the tasks from both Basecamp and TeamGantt. It is just a little option that helps get things done quickly.

Calendar

For the most part the Calendar is pretty self-explanatory. There are some automatic things it provides, as well as features that can be used. Below is a screenshot.

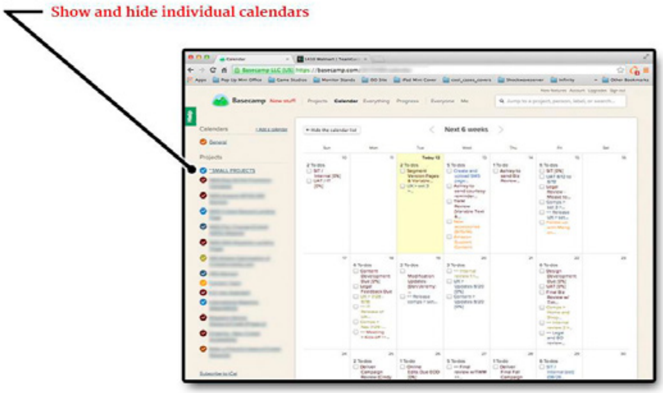

Show and hide individual calendars

To the left of the screenshot, the Calendar shows all the events relating to the projects you are involved with. Be aware that the complete view is on by default, and it may become too much information unless you turn off projects you are not focusing on.

If that is the case, just click the little colored circle with the check mark in order to turn the visibility (on or off) of the related tasks from the calendar. What you see in the larger portion of the calendar are all the to-dos posted by due dates. A user can also use this to checkoff work. Besides the normal To-do items, individuals can also add their own entries and associate them to any accessible project. These, however, do not create a to-do item but, rather, only create a calendar event. Our group uses Google Calendar, so we may not use this as often. However it may be useful to add events if it helps to keep track of events on an individual bases. The following image shows how to create an event. You can add an event by clicking on any of the calendar days. You can add the event's title and additional notes to the desired calendar, and you can even adjust the event to span over multiple days if needed and then email your colleagues.

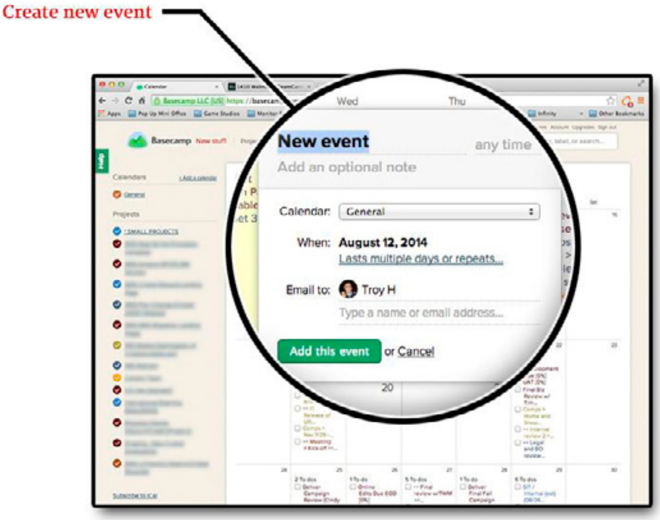

Create new event

New event

any time

Add an optional note

Calendar: General

When: August 12, 2014

Lasts multiple days or repeats...

Email to: Troy H

Type a name or email address...

Add this event or Cancel

Everything

"Everything" is an easy way to browse all items in their respective groups. This section has *Browse every discussion, Review all open to-dos, See every single file, Read all text documents, Show all forwarded emails* and *See all deleted items*.

These are just other ways to find information quickly.

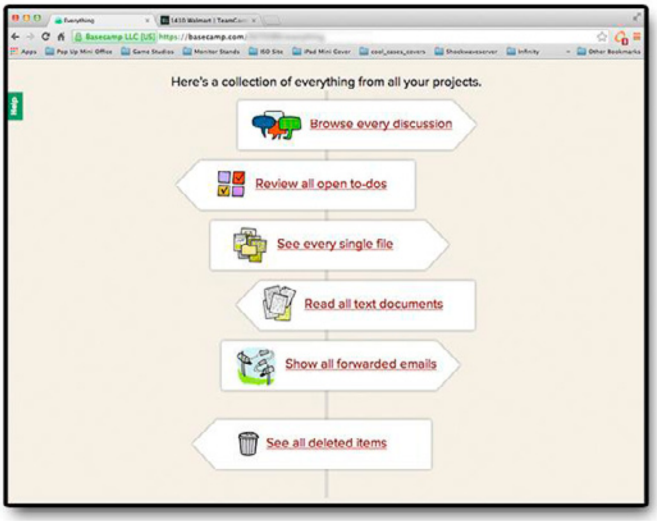

Here's a collection of everything from all your projects.

Browse every discussion

Review all open to-dos

See every single file

Read all text documents

Show all forwarded emails

See all deleted items

To provide a quick breakdown:

- *Browse every discussion* > Provides a listing of any/all textual updates in the order they were added. You can click to get to that discussion and associated project by selecting it.

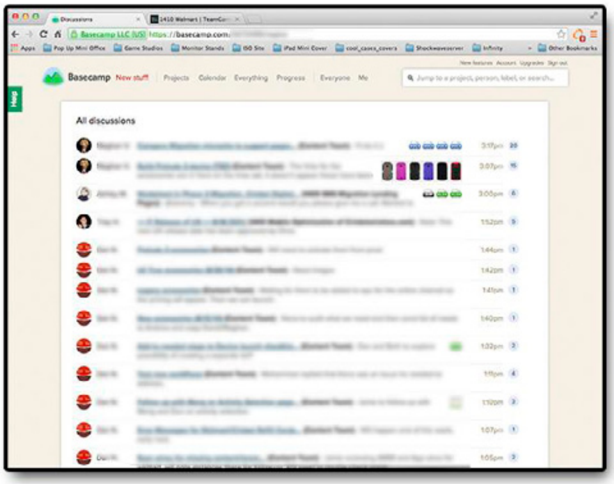

All discussions

- *Review all open to-dos* > Provides access to all the to-dos (that you have access to) that have not been checked off. Again, you can just link to the exact to-do within the project by clicking on any of the items on the page.

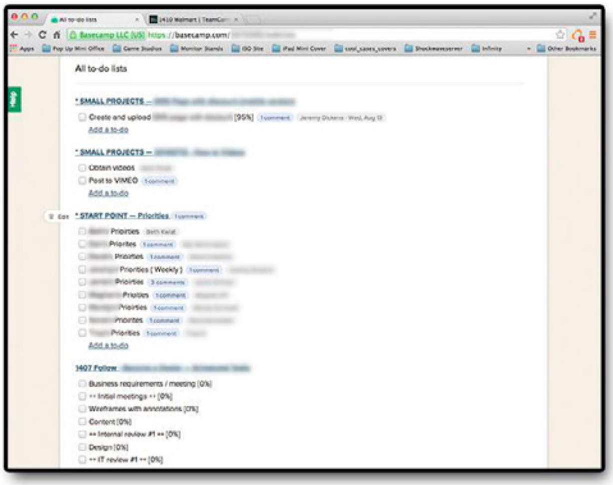

All to-do lists

*SMALL PROJECTS –

□ Create and upload [95%] 1 comment

Add a to-do

*SMALL PROJECTS –

□ Obtain videos

□ Post to VIMEO 1 comment

Add a to-do

*START POINT – Priorities 1 comment

□ Priorities

□ Priorites 1 comment

□ Priorities 1 comment

□ Priorities (Weekly) 1 comment

□ Priorities 3 comments

□ Prioties 1 comment

□ Priorties 1 comment

□ Priorites 1 comment

□ Priorities 1 comment

Add a to-do

1407 Follow

□ Business requirements / meeting [0%]

□ ++ Initial meetings ++ [0%]

□ Wireframes with annotations [0%]

□ Content [0%]

□ ++ Internal review #1 ++ [0%]

□ Design [0%]

□ ++ IT review #1 ++ [0%]

- *See every single file* > Provides access to all uploaded files for all the projects to which you have access to. Useful if you have a lot of projects and you want to have an overview of all uploaded files, etc.
- *Read all text documents* > Shows all documents based on the last update.
- *Show all forwarded emails* > For emails that have responses from outside of Basecamp. We probably will not have a need for this.
- *See all deleted items* > Anything that has been deleted. This is also not used very often.

Progress

Shows who did what in the order it happened. This comes in very handy when wanting to find out any activity of the last few days. Beyond that it may present too much info. The *Progress* section also gives a good indication on who is using Basecamp and how.

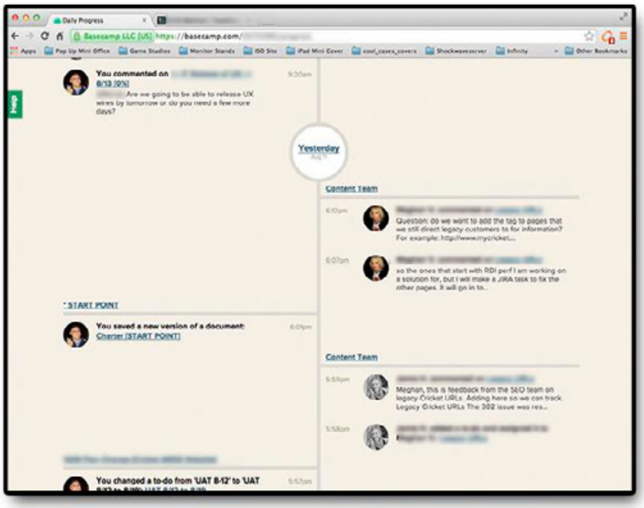

You commented on 8/13 [0%] 9:20am

Are we going to be able to release UX wires by tomorrow or do you need a few more days?

Yesterday

Content Team

6:12pm Question: do we want to add the tag to pages that we still direct legacy customers to for information? For example: http://www.mycricket...

6:07pm so the ones that start with RDI perf I am working on a solution for, but I will make a JIRA task to fix the other pages. It will go in to...

*START POINT

You saved a new version of a document: Charter [START POINT] 6:01pm

Content Team

5:59pm Meghan, this is feedback from the SEO team on legacy Cricket URLs. Adding here so we can track. Legacy Cricket URLs The 302 issue was res...

5:58pm

You changed a to-do from 'UAT 8-12' to 'UAT 5:57pm

You can scroll down, review the messages, files and happenings in real time.

Everyone

This section shows everyone based on the last active individuals. Latest active individuals are posted first.

You can see everyone by clicking the "See all people" link on the bottom left of the screen.

Incidentally, "admins" can add additional people, change access permissions and perform other administrative functions.

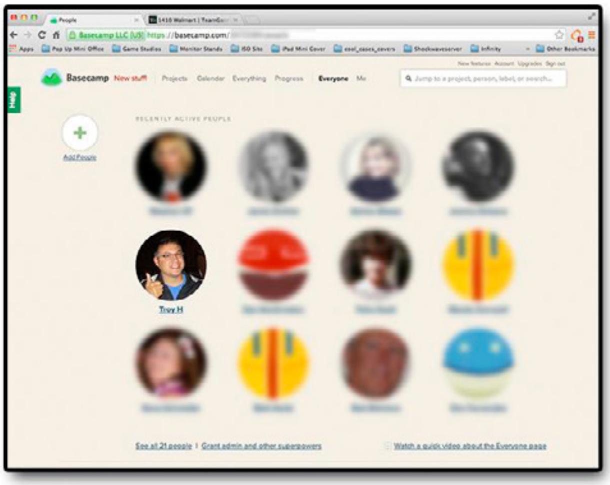

Basecamp New stuff Projects Calendar Everything Progress Everyone Me

RECENTLY ACTIVE PEOPLE

Add People

Troy H

See all 21 people | Grant admin and other superpowers

Watch a quick video about the Everyone page

Me

The "Me" section can be very helpful to quickly see everything on your plate. Provides access to all the *latest activity* across all your projects, *all your open to-dos, recently completed todos* and *files you have shared*.

This should actually be the first place you should go in the morning to see if it lines up with what you know needs to happen.

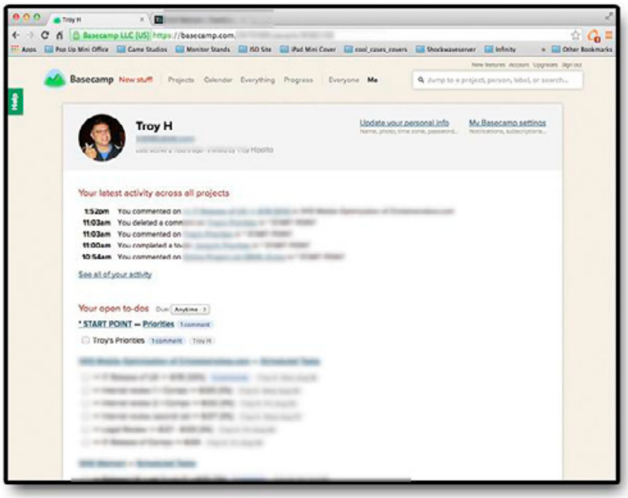

Basecamp New stuff Projects Calendar Everything Progress Everyone Me

Troy H

Update your personal info

My Basecamp settings

Your latest activity across all projects

1:52pm You commented on

11:03am You deleted a comm

11:03am You commented on

11:00am You completed a to-

10:54am You commented on

See all of your activity

Your open to-dos Due: Anytime

*START POINT – Priorities 1 comment

□ Troy's Priorities 1 comment Troy H

How We Use Basecamp (Usage Guidelines)

This is so important we made it a major section in the article. Proper etiquette comes into play when we think of how our actions affect the team and timelines. Basecamp is pretty much an open system. We can use it the way we like. We as a group need to form and follow a sort of protocol or "etiquette" which will be helpful in making everything to start making sense and become more of a natural process.

This process should include "when" and "how" we use Basecamp communications, as well as where information should be located. Basecamp is not a perfect tool by far. It is up to everyone to use and tweak it as needed. At the same time, if we do not report to it, then the information will not be available for everyone else. An added advantage of properly and effectively using the system is that it will actually help reduce the need for some of the meetings and allow you to complete your work.

The rules of today maybe switched later for something that makes more sense. But for now these are the general usage guidelines.

Identifying Your Tasks

Tasks can be anything including replying to messages, reviewing documents, identifying dates and, yes, especially to-dos.

One of the things we have to keep in mind is to look out for each other. If you notice that there are tasks missing that will prevent you from doing your work (or dates that do not seem appropriate) then please let everyone else know. It probably has to be addressed.

So, where do we start?

- *First go to the "Me" menu link*. There you can see what tasks are assigned to you.

  This does not give you a clear priority but it does show you all things you are associated with.

  If you know there are tasks or projects that you have to do work on and it is not there, find out why. If it is not on Basecamp then others may not know it exists. Basecamp should be used as transparently as possible so others can quickly see how the project is going without the need for much interaction.

  Basecamp is designed to show you and others where you are at in the process of your projects.

- Secondly (and this is optional): *START POINT* project is a good place to go to, as there is a to-do list called *Priorities*. There you can find your name and make notes on your goals for the day *in order of importance*. Creating this priority list helps you focus on items needed. And we all know the focus can change at any moment. So feel free to update that to-do item for any updates.
- Once you are focused for the day, it is best to dive in the individual project you are working on. Based on the priority list you have created, then go the project you have to work on.

Identifying project scope and important files and links

Even after you have reviewed the project and your to-do's, do not overlook the possibility that there are times when there may be context missing on the project or you may need access to some particular bit of additional information. This could be a reference info, a "what the heck is this project about," a list of who is on it, or information on how to get access to the needed file(s).

It may just be that you need more details on the actual to-do/task...

Formatting explanation of the to-dos schedule

To make the to-dos a little more precise and at the same time keep the amount of content readable for the to-dos, we have implemented a subject formatting technique to assist in this matter.

As we mentioned before, the to-dos are tasks that can be also arranged as a schedule (view section 1.3 for general details on to-do's). This subsection, however, is really designed to break down why we format it the way we do.

Please note that not all to-do items may be formatted this way. However, if you have a *series of tasks* that form a schedule, then it is best to use these practices. To simplify the different styles let's call these *series of tasks* (to-do schedules) and *non-series of tasks* (one-offs or similar tasks).

To illustrate the point the images that follow are 2 different views of the same *series of tasks*. The first image is what you see from Basecamp, while the second is a more visual timeline generated in TeamGantt.

Series of tasks:

Basecamp view. Most projects that require a series of tasks are broken up in usable chunks. They tend to include most of the larger events, but often meetings spring up for additional reviews or issues that are not related to the project.

You can see this project is a little more complicated than most, but the basic structure is pretty standard. Typically, most projects have:

- BO (Business Owner) ZIP and CB (Creative Brief). In the example below it was already completed and checked off.
- Internal kick-off. Again, this already happened and has been checked off.
- UX development
- Content support
- Comps
- Internal reviews (some times these are not listed as the dates shift too often)
- Legal reviews
- IT Release
- UAT prep
- IT Dev
- UAT internal testing
- Launch

As you will notice, the following image has multiple releases. Sometimes this is needed if you are dividing design and development/IT groups. Dev/IT may have to get started on a project overlapping the design schedule in order to make launch dates.

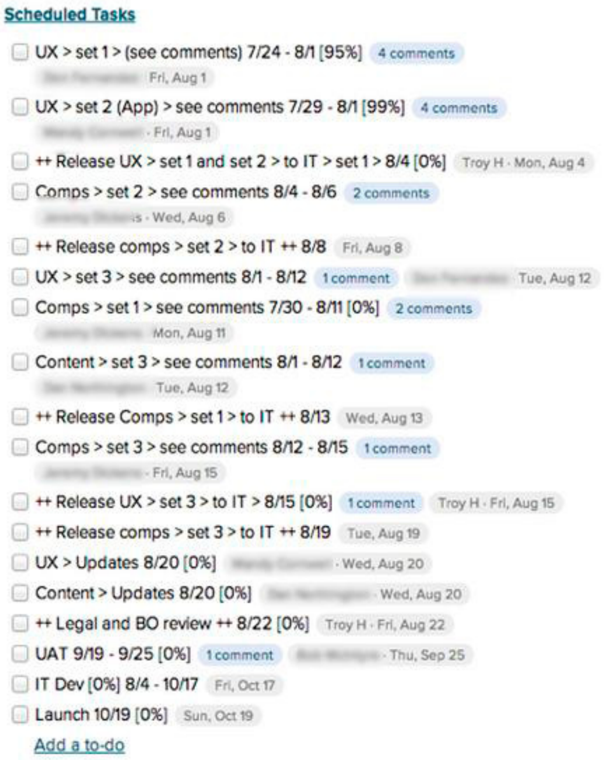

Scheduled Tasks

□ UX > set 1 > (see comments) 7/24 - 8/1 [95%] 4 comments
[redacted] Fri, Aug 1
□ UX > set 2 (App) > see comments 7/29 - 8/1 [99%] 4 comments
[redacted] · Fri, Aug 1
□ ++ Release UX > set 1 and set 2 > to IT > set 1 > 8/4 [0%] Troy H · Mon, Aug 4
□ Comps > set 2 > see comments 8/4 - 8/6 2 comments
[redacted]s · Wed, Aug 6
□ ++ Release comps > set 2 > to IT ++ 8/8 Fri, Aug 8
□ UX > set 3 > see comments 8/1 - 8/12 1 comment [redacted] Tue, Aug 12
□ Comps > set 1 > see comments 7/30 - 8/11 [0%] 2 comments
[redacted] Mon, Aug 11
□ Content > set 3 > see comments 8/1 - 8/12 1 comment
[redacted] Tue, Aug 12
□ ++ Release Comps > set 1 > to IT ++ 8/13 Wed, Aug 13
□ Comps > set 3 > see comments 8/12 - 8/15 1 comment
[redacted] · Fri, Aug 15
□ ++ Release UX > set 3 > to IT > 8/15 [0%] 1 comment Troy H · Fri, Aug 15
□ ++ Release comps > set 3 > to IT ++ 8/19 Tue, Aug 19
□ UX > Updates 8/20 [0%] [redacted] · Wed, Aug 20
□ Content > Updates 8/20 [0%] [redacted] · Wed, Aug 20
□ ++ Legal and BO review ++ 8/22 [0%] Troy H · Fri, Aug 22
□ UAT 9/19 - 9/25 [0%] 1 comment [redacted] · Thu, Sep 25
□ IT Dev [0%] 8/4 - 10/17 Fri, Oct 17
□ Launch 10/19 [0%] Sun, Oct 19
Add a to-do

Now that we have an idea of the different groupings, let's consider formatting. If you take another look at the series of tasks shown above, you will notice that some have "++" in front of them and others do not.

- *The "++" prefix* generally represents a *release* or *major meeting*. These are what we call non-tasks or things that do not require the online group to develop. Items *without* the "++" prefix generally represent design of UX, content creation or comp designing.
- You will also notice that we use the *overarching subject first*. So you may read things like *UX*, *Content* or *Comps*. Then you will notice a little arrow like this ">". After that a small amount of detail (just enough to understand what is being worked on).
- After the detail you will often see a *date* or *date range*. For events that take one day, a simple date is needed, while for events spanning a period of time it is good to just put the date in the description. This is important because Basecamp only tracks end dates. We want to show the start and end dates.
- After the date or date range you may see a *percentage in brackets* like this – [50%]. This is manually filled out so we get an idea on where this task is in the process. It is also auto-translated visually in the Gantt chart in TeamGantt. And if the description is a bit vague, that is why you add a comment to it. That way anyone looking at this particular to-do can see that there is a comment which can be clicked on to drill down and see additional details.
- Then the task is *assigned to someone* and given a *due date*. You may be curious why there is a date range and a due date. For starters, while due dates are tracked in the system, we format these dates visually to make it easier to read and identify start dates.

Also TeamGantt has start dates as well as end dates. It is easier to adjust the start in TeamGantt once you can actually read the date ranges in the description.

So let's use this example and break it down.

```
Comps > set 1 > see comments > 8/12 – 8/15 [0%] 1 comment
  Person Name Mon Aug 11
```

- `Comps` = Overall subject.
- `set 1` = Short detail.
- `see comments` = A note signifying that more details on the tasks are to found in the comments.
- `8/12 – 8/15` = Start and end date (note we will probably have several internal approval meetings in between these dates that may or may not be notated on the schedule).
- `[0%]` = The estimated percentage of the task completion.
- `1 comment` = Shows how many comments are associated with this task.
- `Person Name Mon Aug 11` = Who is assigned (and defaults emails to), and when is that task due.

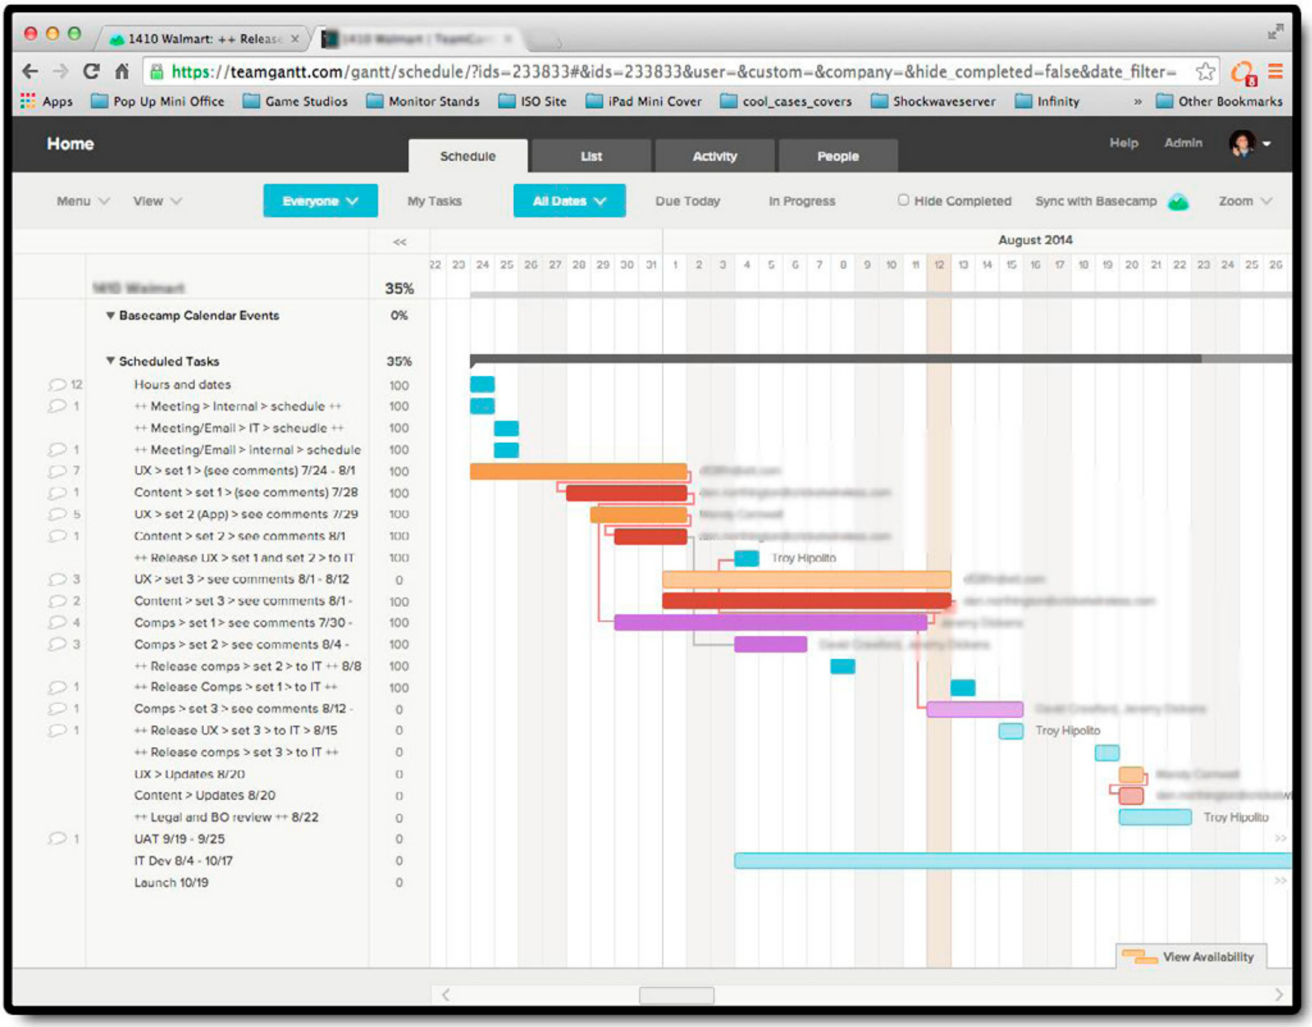

1410 Walmart: ++ Release

https://teamgantt.com/gantt/schedule/?ids=233833#&ids=233833&user=&custom=&company=&hide_completed=false&date_filter=

	[redacted]	35%
	▼ Basecamp Calendar Events	0%
	▼ Scheduled Tasks	35%
12	Hours and dates	100
1	++ Meeting > Internal > schedule ++	100
	++ Meeting/Email > IT > scheudle ++	100
1	++ Meeting/Email > internal > schedule	100
7	UX > set 1 > (see comments) 7/24 - 8/1	100
1	Content > set 1 > (see comments) 7/28	100
5	UX > set 2 (App) > see comments 7/29	100
1	Content > set 2 > see comments 8/1	100
	++ Release UX > set 1 and set 2 > to IT	100
3	UX > set 3 > see comments 8/1 - 8/12	0
2	Content > set 3 > see comments 8/1 -	100
4	Comps > set 1 > see comments 7/30 -	100
3	Comps > set 2 > see comments 8/4 -	100
	++ Release comps > set 2 > to IT ++ 8/8	100
1	++ Release Comps > set 1 > to IT ++	100
1	Comps > set 3 > see comments 8/12 -	0
1	++ Release UX > set 3 > to IT > 8/15	0
	++ Release comps > set 3 > to IT ++	0
	UX > Updates 8/20	0
	Content > Updates 8/20	0
	++ Legal and BO review ++ 8/22	0
1	UAT 9/19 - 9/25	0
	IT Dev 8/4 - 10/17	0
	Launch 10/19	0

TeamGantt project is just pulling the Basecamp information, providing a visual reference of the time it takes to complete a task (start and end date), a visual percentage of completion, and some of the same tools that Basecamp has.

To simplify, the Project Managers generally set up the permissions for TeamGantt to be "view only." This is so individuals do not have to try to adjust things from there. But the Project Managers do have the ability to create messages and tasks from there if they choose (or need) to.

The Project Managers also try to color-code 4 different types of tasks: Non-tasks (the ones with the "++" prefix): a default powder blue:

- UX: light orange
- Content: orange-red
- Comps: a hot pink/fuchsia color.

In the previous graphic you will notice that a few UX tasks are 95% done. The timeline items for that tasks is actually 95% full. This is a visual indicator of where the task is.

Non series set of tasks:

The following image is example of a *non series set of tasks*. These are things like one-offs or recurring tasks.

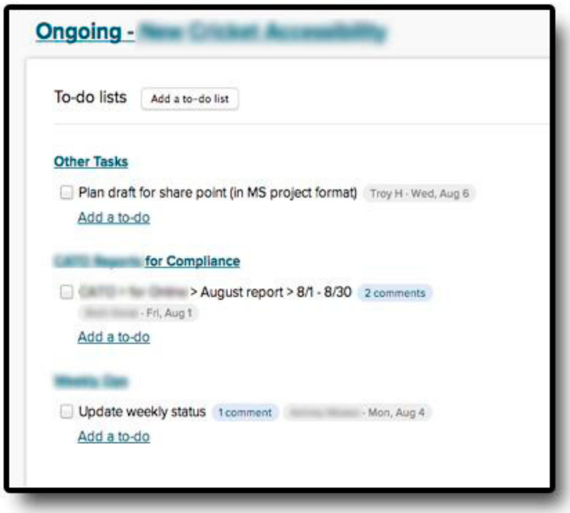

Ongoing - [redacted]

To-do lists Add a to-do list

Other Tasks

□ Plan draft for share point (in MS project format) Troy H - Wed, Aug 6

Add a to-do

[redacted] for Compliance

□ [redacted] > August report > 8/1 - 8/30 2 comments

[redacted] - Fri, Aug 1

Add a to-do

[redacted]

□ Update weekly status 1 comment [redacted] - Mon, Aug 4

Add a to-do

The non-series of tasks are usually simple, yet explain things when possible in the subject first type of formatting.

When to use regular email vs. Basecamp messages/to-dos

There are lots of messages that do not have to be tracked or which don't specifically pertain to a task. If this is the case, you don't have to use Basecamp. You will have to decide if you want the message to be seen by others or not.

We try to streamline whenever possible but also communicate enough to complete the tasks and "asks" at hand.

These are some examples of what *not to post on Basecamp*. "Hi John – how was your weekend?", "I did not like the meeting and thought it was a bad idea" or "I am concerned I committed to a deadline I can not reach." These are examples of messages or personal conversations that should handled *outside* of Basecamp.

Also if we are emailing other people outside of Basecamp, we should just use an email. They will not know what the email is if it coming from the Basecamp system.

Replying to Basecamp messages/to-dos

One obstacle we tend to run across is that we do not always have an understanding that a Basecamp-generated email is an email that requests a reply. If emails are not acknowledged then it can have an adverse effect when trying to finish projects in a timely manner. *As such:*

- Always direct the email to the main person when applicable.
  For example: If you are sending out a message and have selected several individuals to receive it, please direct the message to the individuals by adding @ symbol followed by the persons' names. It will end up looking something like this: @Jon.
  That way once they receive the email the first thing they will see is who the email focused on.
- Please reply to Basecamp emails.
- Please start conversations on to-dos and messages on Basecamp (when appropriate).

More information about general usage of Basecamp messages and to-dos please view sections 2.2 and 2.3.

Conclusion

Basecamp is a tool to allow us to centralize conversations, help build and maintain task-driven timelines. It can also be integrated with a number of other tools like *TeamGantt*, which allows us to visually see the timelines (beginning to end), percentage of items completed, certain print features and just allows the online group to quickly stay on track.

Although Basecamp is a great tool to have, it only works well when people are consciously using it in a productive manner. This is a flexible system that requires a little manual work to keep things running smoothly. Just remember, you can also take a peek at the online help section where there are guides, videos and cheat sheets at: *https://basecamp.com/help*.

I hope this article gives you the basics to help get projects done a little more efficiently and with more peace of mind. Having a sense of control and being able to confidently get things done and report positively to the client definitely makes your life easier. I like easy (Freudian slip).

There are of course more advanced tools but we have chosen Basecamp and related online applications because they are flexible enough for the projects we are working with, while easy enough for clients to respond to. All things are centralized and documented automatically in one place. And if the client responds, then you have a direction you can move towards to make it closer to finishing the goals of the job.

If anyone is interested in the work we have done please take a gander at our site *http://www.isointeractive.com* as we have at least some of our public projects posted there. Mostly we deal with helping clients and partners fixing or developing mobile apps, websites, software reviews/audits, games, 3D simulations, lots of specialty projects and good old web development. Typically they range from the range of 10k to under a million USD.

A few links of interest:

- ISO White Paper: *www.isointeractive.com/pdf*
- ISO Video: *www.isointeractive.com/#showreel*
- ISO Website: *www.isointeractive.com*

Thank you and we look forward to continue contributing to the interactive community.

If you have any needs or even just want to brainstorm, please feel free to connect.

- email: *troy@isointeractive.com*
- skype: *troyhipolito*
- web: *isointeractive.com*
- facebook: *facebook.com/ISOinteractive*
- twitter: *@isointeractive*
- instagram: *iso_interactive*

ABOUT THE AUTHOR

Troy Hipolito is the Senior Consultant at ISO Interactive (a consulting social and mobile game company that supports agencies for campaigns, Facebook games, iPhone Apps and that sort of thing).

A Complete Guide to FreeNAS Hardware Design,

Part III: Pools, Performance, and Cache

JOSHUA PAETZEL

ZFS Pool Configuration

ZFS storage pools are comprised of vdevs which are striped together. vdevs can be single disks, N-way mirrors, RAIDZ (Similar to RAID5), RAIDZ2 (Similar to RAID6), or RAIDZ3 (there is no hardware RAID analog to this, but it's a triple parity stripe essentially). A key thing to know here is a ZFS vdev gives the IOPs performance of one device in the vdev. That means that if you create a RAIDZ2 of ten drives, it will have the capacity of 8 drives but it will have the IOPs performance of a single drive. The need for IOPs becomes important when providing storage to things like database servers or virtualization platforms. These use cases rarely utilize sequential transfers. In these scenarios, you'll find larger numbers of mirrors or very small RAIDZ groups are appropriate choices. At the other end of the scale, a single user trying to do a sequential read or write will benefit from a larger RAIDZ[1|2|3] vdev. Many home media server applications do quite well with a pool comprising a single 3-8 drive RAIDZ[1|2|3] vdev.

RAIDZ1 gets a special note here. When a RAIDZ1 loses a drive, all the other drives in the vdev become single points of failure. A ZFS storage pool will not operate if a vdev fails. This means if you have a pool made up of a single 10 drive RAIDZ vdev and one drive fails, pool operation depends on none of the remaining 9 drives failing. In addition, with modern drives being as large as they are, rebuild times are not trivial. During the rebuild period, all of the drives are doing increased I/O as the array rebuilds. This additional stress can cause additional drives in the array to fail. Since a degraded RAIDZ1 can withstand no additional failures, you are very close to "game over" there. Powers of 2 pool configuration: there is much wisdom out there on the internet about the value of configuring ZFS vdevs in a power of two. This made some sense when building ZFS pools that did not utilize compression. Since FreeNAS utilizes compression by default (and there are 0 cases

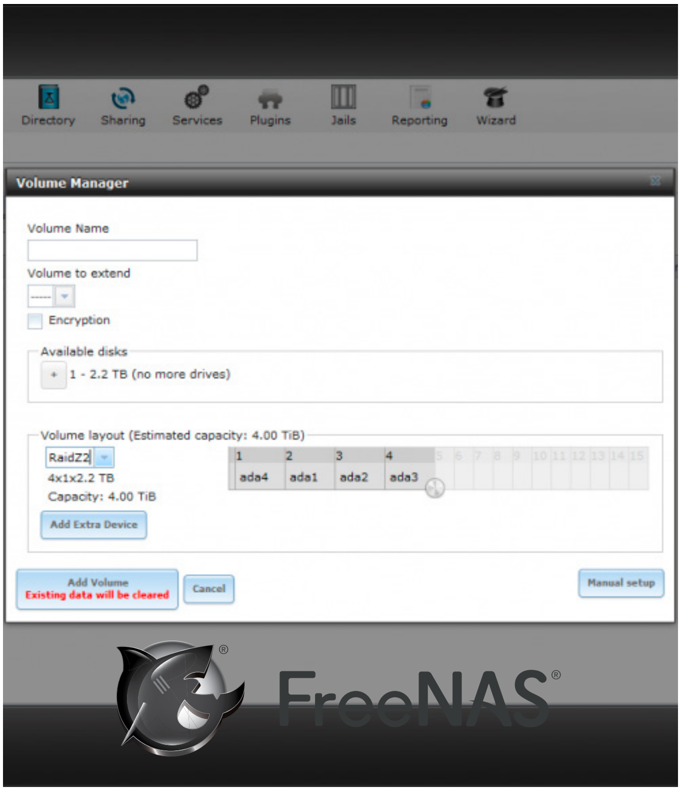

Volume Name

Volume to extend: ------

□ Encryption

Available disks: 1 - 2.2 TB (no more drives)

Volume layout (Estimated capacity: 4.00 TiB)

RaidZ2
4x1x2.2 TB
Capacity: 4.00 TiB

1	2	3	4
ada4	ada1	ada2	ada3

where it makes sense to change the default!), any attempts to optimize ZFS with the vdev configuration are foiled by the compressor. Pick your vdev configuration based on the IOPs needed, space required, and desired resilience. In most cases, your performance will be limited by your networking anyway.

ZIL Devices

ZFS can use dedicated devices for its ZIL (ZFS intent log). This is essentially the write cache for synchronous writes. Some workflows generate very little traffic that would benefit from a dedicated ZIL, others use synchronous writes exclusively and, for all practical purposes, require a dedicated ZIL device. The key thing to remember here is the ZIL always exists in memory. If you have a dedicated device, the memory ZIL is mirrored to the dedicated device, otherwise it is mirrored to your pool. By using an SSD, you reduce latency and contention by not utilizing your data pool (which is presumably comprised of spinning disks) for mirroring the in-memory ZIL. There's a lot of confusion surrounding ZFS and ZIL device failure. When ZFS was first released, dedicated ZIL devices were essential to data pool integrity. A missing ZIL vdev would render the entire pool unusable. With these older versions of ZFS, mirroring the ZIL devices was essential to prevent a failed ZIL device from destroying the entire pool. This is no longer the case with ZFS. Missing ZIL vdevs will impact performance but will not cause the entire pool to become unavailable. However, the conventional wisdom that the ZIL must be mirrored to prevent data loss in the case of ZIL failure lives on. Keep in mind that the dedicated ZIL device is merely mirroring the real in-memory ZIL. Data loss can only occur if your dedicated ZIL device fails and the system crashes with writes in transit in the unmirrored memory ZIL. As soon as the dedicated ZIL device fails, the mirror of the in-memory ZIL moves to the pool (in practice, this means you have a window of a few seconds where a system is vulnerable to data loss following a ZIL device failure). After a crash, ZFS will attempt to replay the ZIL contents. SSDs themselves have a volatile write cache, so they may lose data during a bad shutdown. To ensure the ZFS write cache replay has all of your inflight writes, the SSD devices used for dedicated ZIL devices should have power protection. HGST makes a number of devices that are specifically targeted as dedicated ZFS ZIL devices. Other manufacturers such as Intel offer appropriate devices as well. In practice, only the designer of the system can determine if the use case warrants a professional enterprise grade SSD with power protection or if a consumer-level device will suffice. The primary characteristics here are low latency, high random write performance, high write endurance, and, depending on the situation, power protection.

L2ARC Devices

ZFS allows you to equip your system with dedicated read cache devices. Typically, you'll want these devices to be lower latency than your main storage pool. Remember that the primary read cache used by the system is system RAM, which is orders of magnitude faster than any SSD. If you can satisfy your read cache requirements with RAM, you'll enjoy better performance than if you use SSD read cache. In addition, there is a scenario where an L2ARC read cache can actually drop performance. Consider a system with 6GB of memory cache (ARC) and a working set that is 5.9 GB. This system might enjoy a read cache hit ratio of nearly 100%. If SSD L2ARC is added to the system, the L2ARC requires space in RAM to map its address space. This space will come at the cost of evicting data from memory and placing it in the L2ARC. The ARC hit rate will drop, and misses will be satisfied from the (far slower) SSD L2ARC. In short, not every system can benefit from an L2ARC. FreeNAS includes tools in the GUI and at the command line that can determine ARC sizing and hit rates. If the ARC size is hitting the maximum allowed by RAM, and if the hit rate is below 90%, the system can benefit from L2ARC. If the ARC is smaller than RAM or if the hit rate is 99.X%, adding L2ARC to the system will not improve performance. As far as selecting appropriate devices for L2ARC, they should be biased towards random read performance. The data on them is not persistent, and ZFS behaves quite well when faced with L2ARC device failure. There is no need or provision to mirror or otherwise make L2ARC devices redundant, nor is there a need for power protection on these devices.

ABOUT THE AUTHOR

iXsystems Director of IT

Does your Information Belong to the CIA Triad?

ROB SOMERVILLE

Confidentiality, Integrity and Availability are the three pillars of Information Security. In this article, we pose a number of scenarios to you, the IT professional, and ask "What would you do"? Every environment is different, so we will not provide any answers. Rather, we want to stimulate thought and debate around the ethics that Donn Parker says are missing from the computer center.

01

Question 1.

The IT help-desk is understaffed, and you are the only member of IT available. A new member of staff requests a password change, but this is a help-desk call and they are hot-desking. Do you change it anyway or how do you proceed?

02

Question 2.

You are the web-master of a large corporate site, and the senior manager responsible for adding content for his department is on leave. A major crisis has developed around a particular issue and a senior manager asks you to publish a revised Word document concerning the matter. You discover this document still has all the modifications and revisions highlighted and available, and looks very unprofessional. Also, you suspect that anyone reading the document will glean information that is not desired. Do you mention this to the manager?

03

Question 3.

You politely point this out to the manager verbally but are reprimanded and told to mind your own business. What now?

04

Question 4.

What steps do you take when disposing of large quantities of end of life hardware and consumables? Securely wiping each device takes too much time, and there is no budget for a third party to provide this service. How do you handle difficult items like cartridge ribbons that retain an imprint of printouts and photocopiers that have data stored in internal hard drives?

05

Question 5.

When designing a database application, what criteria do you use as to when data is encrypted? Should this be performed at the application or database layer and what implications does this have for disaster recovery? When is database encryption useless?

06

Question 6.

When designing and commissioning a project, when is it best practice to place the software in escrow? If a vendor refuses to permit this, yet you are under pressure from management to commit, what action can you take?

07

Question 7.

How often do you test your backups? Do you just sample the data or do you perform "Restore from bare metal" tests? If your organisation was to suffer fire, flood or earthquake, are there up to date copies off-site?

08

Question 8.

A critical project requires multiple leased lines for redundancy. Do you use the same vendor for both? Do you need to ensure you have replacement pre-configured "Like for Like" routers etc. available?

09

Question 9.

An open Wi-Fi hotspot with a strong signal is close to your premises. Connection to this node is not password protected. What risk, if any, does this pose to your Internet connected LAN if staff members with laptops connect to your LAN via cable and the Internet via external Wi-Fi? What steps can you take to mitigate any risk?

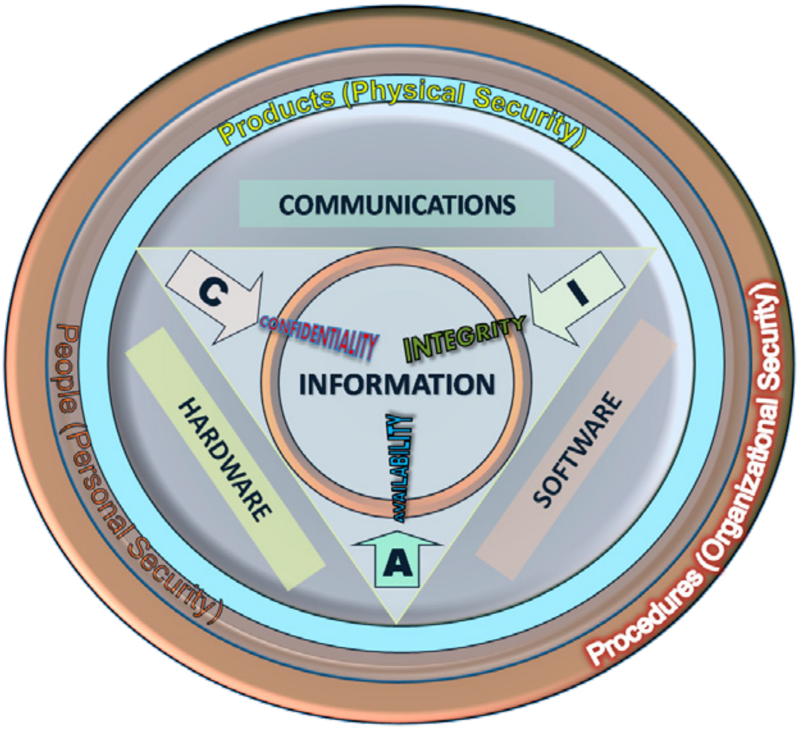
Products (Physical Security)
COMMUNICATIONS
C
I
CONFIDENTIALITY
INTEGRITY
INFORMATION
AVAILABILITY
A
HARDWARE
SOFTWARE
People (Personal Security)
Procedures (Organizational Security)

formation belong to the CIA triad?
Image courtesy of John M. Kennedy T.

10

Question 10.

A laptop with encrypted data is infected with a virus in the form of a root-kit. You do not have access to the keys to mount the drive from a separate boot disk but you do have the administrator password for the machine. How do you gain access to the hard disk to disable the malware from loading as it loads prior to the O/S fully loading?

11

Question 11.

Does your organisation have a social media policy about the use of Facebook and Twitter etc. during working hours? Outside business hours? Does this include personal devices? If so does it cover both personal and postings in an official capacity?

12

Question 12.

What steps does your organisation take to prevent misrouted email data loss e.g. by picking the wrong "Smith" from a distribution list? Are documents pro-actively marked (e.g. confidential, for general release etc.)? How easy is it to spoof a user in your organisation if you telnet to port 25 of your email server?

13

Question 13.

What is the minimum level of password complexity demanded in your organisation? Do you use single sign-on? If the password levels are complex, do staff write the passwords on Post-It notes and place them on their monitor etc? How often are password changes enforced? What is the major downside of single sign-on systems?

14

Question 14.

Do you have access via a separate route to the Internet that is not firewalled or connected to your corporate LAN? What could this be used for in diagnosing a major systems outage e.g. email or web-server? Do you use this for accurate penetration testing?

15

Question 15.

The plastic soft-touch keypad on the door-lock of the Data-centre has 4 discoloured digits due to heavy use over the years. Assuming no numbers are repeated in the entry code, how many permutations would it take to brute-force the combination? Does this warrant replacing the keypad? Will the lock fail if the combination is entered incorrectly too many times? How often is the password changed?

ABOUT THE AUTHOR

Rob Somerville has been passionate about technology since his early teens. A keen advocate of open systems since the mid-eighties, he has worked in many corporate sectors including finance, automotive, airlines, government and media in a variety of roles from technical support, system administrator, developer, systems integrator and IT manager. He has moved on from CP/M and nixie tubes but keeps a soldering iron handy just in case.

Great Specials

On FreeBSD® & PC-BSD® Merchandise

Give us a call & ask about our
SOFTWARE BUNDLES
1.925.240.6652

$39.95
FreeBSD 9.1 Jewel Case CD Set
or FreeBSD 9.1 DVD

$29.95
PC-BSD 9.1 DVD

$49.95
Save a BUNDLE!
The PC-BSD 9.0 Users Handbook
PC-BSD 9.1 DVD

$99.95
The FreeBSD CD **or** DVD Bundle

Inside each CD/DVD Bundle, you'll find:
FreeBSD Handbook, 3rd Edition
Users Guide FreeBSD Handbook, 3rd Edition, Admin Guide
FreeBSD 9.1 CD or DVD set
FreeBSD Toolkit DVD

Stylish Dress Attire
Look Your Professional Best

Comfy Apparel
Stay Warm in Zip Ups & Pullovers

T-Shirts
Lots of Styles to Choose From

FreeBSD 9.1 Jewel Case CD/DVD	$39.95

CD Set Contains:

Disc 1	Installation Boot LiveCD (i386)
Disc 2	Essential Packages Xorg (i386)
Disc 3	Essential Packages, GNOME2 (i386)
Disc 4	Essential Packages (i386)

FreeBSD 9.0 CD	$39.95
FreeBSD 9.0 DVD	$39.95

FreeBSD Subscriptions

Save time and $$$ by subscribing to regular updates of FreeBSD

FreeBSD Subscription, start with CD 9.1	$29.95
FreeBSD Subscription, start with DVD 9.1	$29.95
FreeBSD Subscription, start with CD 9.0	$29.95
FreeBSD Subscription, start with DVD 9.0	$29.95

PC-BSD 9.1 DVD (Isotope Edition)

PC-BSD 9.1 DVD	$29.95
PC-BSD Subscription	$19.95

The FreeBSD Handbook

The FreeBSD Handbook, Volume 1 (User Guide)	$39.95
The FreeBSD Handbook, Volume 2 (Admin Guide)	$39.95

The FreeBSD Handbook Specials

The FreeBSD Handbook, Volume 2 (Both Volumes)	$59.95
The FreeBSD Handbook, Both Volumes & FreeBSD 9.1	$79.95

PC-BSD 9.0 Users Handbook	$24.95
BSD Magazine	$11.99
The FreeBSD Toolkit DVD	$39.95
FreeBSD Mousepad	$10.00
FreeBSD & PCBSD Caps	$20.00
BSD Daemon Horns	$2.00

Bundle Specials!
Save $$$

Just Plain Fun
Mousepads & Novelty Horns

BSD Magazine
Available Monthly

FreeBSD Mall

For even MORE items
visit our website today!
www.FreeBSDMall.com

Google Earth Forensics

Using Google Earth Geo-Location in Digital Forensic Investigations

Digital Forensics 101

MICHAEL HARRINGTON, MICHAEL CROSS

Digital Forensics is a branch of forensic science that focuses on the recovery, examination, and investigation of evidence stored on computers and other digital devices, as well as various media that may have been used to store data. Although it is commonly associated with criminal investigations, digital forensics has been used in civil cases, internal investigations, tribunals, and other inquiries or forums that require an exploration of data.

The process of performing a digital forensic investigation can be broken down into four stages:

- Seizure, in which computers, mobile devices and other devices and/or media are obtained and preserved.
- Acquisition, in which the data is retrieved from a device
- Analysis, in which an image or copy of the data acquired in the previous step is examined
- Reporting, in which the procedures and processed that were followed in the previous steps are documented, along with the evidentiary findings

Seizure

When a computer or other device is seized, it is taken into custody and secured with goal of preserving any potential evidence. As with every stage of a digital forensic investigation, you will document the scene, actions that were taken, and procedures that were followed. It is also important at this stage to establish a chain of custody that will carry on through all the other stages, documenting who and when and when a person had position of evidence.

In addition to photographing the scene where the computer or device was seized, photograph the computer or mobile device and what is displayed on the screen. Photographing the screen will preserve what applications were open, possible information, and will show what the user was last using doing on the computer or device. Under no circumstances should you use the computer/device, search for evidence, or alter its running condition. A rule of thumb is that if it is turned off, leave it off; if it is turned on, leave it on.

During the seizure, some steps may be taken to acquire digital evidence. If a computer is turned on, you would start by collecting any live data, inclusive to taking an image of the physical memory. A utility that can be used to image the RAM is F-Response (*www.f-response.com*). This tool could also be used to collect a logical image of the disk if you discovered the hard disk was encrypted.

You would also gather any other data that is required for the investigation about the computer's live state, such as logged on users, its network connection state, running processes, and so on.

You should also take effort in documenting how the computer or device was found. Photographs and diagrams should be made of how it was setup when found, inclusive to any cords plugged into the machine. You should also label all of the cords, and document the model numbers and serial numbers of the computer/device and any other devices attached to it. Nothing should be disconnected from a computer or device until the previous steps have been completed.

When you are ready to transport the computer/device, you should package all of the components in anti-static bags, and seize any other storage media. This would include external hard disks, USB sticks, as well as CDs and DVDs that may contain data. To keep the media safe, you should avoid putting it near anything that may damage the data, such as magnets, radio transmitters, and so on. In gathering these additional items, you should also collect any manuals or documentation that may be related to the device. You never know if these will be helpful later in your investigation, or if they contain useful information (such as passwords, etc.).

There are additional considerations when a mobile device is seized. When a mobile device is connected to a cellular network, it may access new data that will overwrite evidence. Similarly, a mobile GPS unit that is turned on mayncontinue to record track points (i.e., locations that the GPS has been) as its being transported. Because a mobile phone or tablet can be sent a command to wipe the device, you also run the risk of everything on it being erased. To preserve potential evidence on a mobile phone, GPS or other device, it is important they are stored in a Faraday bag or cage. A Faraday cage is an area protected by material that blocks signals, essentially creating the same conditions of being in a "dead zone" where you cannot get a cell phone signal from your carrier. A Faraday bag is used to store mobile devices for transport, preserving any evidence stored on them.

Acquisition

The acquisition stage is where data is retrieved from a device or media, and generally occurs after the evidence has been collected, safeguarded and transported. In acquiring evidence from a device, a decision is made whether you need to perform a live or dead analysis. A live analysis is performed when a computer or device is powered on, and cannot be powered off until this information is collected. A dead analysis occurs when the machine is powered off, and transported to a lab where data can be retrieved in a controlled environment.

Acquiring data from a computer, device, or various media that may be used to store potential evidence generally requires specialized tools. This is not to say there are not times when a mobile device may require the manual acquisition of data, whereby an investigator uses the user interface of a phone or other device to view and photograph information displayed on the screen. However, in doing so, the only data that will be displayed is that which is accessible to the device's operating system and/or apps. In addition, using the interface may result in data being written to the device. To safely acquire all of the data, inclusive to that which may have been deleted, software and hardware tools are commonly used to create a bit-for-bit copy of what is stored on the device. Once a copy of the data is acquired, the investigator can then examine the copy of the data so that the original remains untouched during analysis.

There are several ways in which you may acquire a copy of what is stored on a file system, but not all of them will provide the same results. These methods include:

- Copying files, which will only copy the files that are on the system and not ones that may have been deleted. Also, metadata related to file ownership, times a file was accessed, permissions and other data may be lost in copying the file.
- Backups, which will restore a copy of the files. Depending on the backup software used, not all of the metadata related to files will be included with the backup, and it will not capture information about deleted files.
- Copying disk partitions, which will create a bit-for-bit copy of the file system including metadata related to the files and information residing in unallocated space.
- Copying the entire disk, which creates a bit-for-bit copy of the file system, including storage space before and after disk partitions.

In looking at these methods, you can see that a bit-by-bit copy of the data will yield the most possible results. While you might think this would only apply to the hard disk of a computer, many mobile devices use file systems and may be used as storage devices. In addition, devices that use SD cards can have the card removed and processed like other removable media. By using various tools discussed later in this chapter, you will be able to collect the data on these devices, making a copy that you can then analyze to identify evidence related to your case.

Analysis

The analysis stage generally occurs after evidence has been collected. If live data is not being examined, then an investigation is conducted against static data that has been copied from a system. Once an image of data on the computer, device, or other media has been made, an examination of the data takes place. This may involve performing keyword searches relating to a crime, running scripts to identify certain types of data, manually reviewing information and content of files, and various other techniques.

By analyzing various types of data found on a machine, investigators will search for evidence that implicates or exonerates a suspect. The evidence may include digital photographs or downloaded images (as in the case of child pornography cases), electronic spreadsheets (in the case of financial crimes), email and other types of data. Using the content, metadata, or other information discovered, the investigator may reconstruct a series of events related to the case.

Reporting

Documentation is crucial to any digital forensics case. It is important to make a record of any actions taken, devices or media examined, procedures that were followed, and other details relating to the evidence. Remember that, especially after a case goes to court, there is the possibility that anything related to the case may be questioned, and your documentation may be used to provide answers.

Throughout the process of conducting an investigation, it is vital that the integrity of the data and the device storing it is preserved, and part of this involves a documented chain of custody. Once a computer, device or media is seized, it should start the chain of custody, showing who initially took possession and who had custody of it after that point. It is also important to remember that the original devices, storage media, or other items that evidence was collected from may be requested by defense council or other parties involved in the case. In some cases, evidence files or images taken of a system may be requested. By preserving these items and ensuring there is a record of who had access to them, you can help to ensure the evidence has not been corrupted or tampered with in anyway.

It should also come as no surprise that you will need to create a report about what was found during the course of your investigation, and how it applies to the case. This could include listings and details about any files found on storage mediums (e.g., hard disks, tape, USB devices, etc.), information recovered from emails or other sources, and any other data that is being used as evidence. As we will discuss later in this chapter, many commercial tools provide features that will automatically generate reports about the files that were found. You would also write a report yourself that outlined the steps taken to acquire and analyze the data, and how the files or information found apply to the case. The reports themselves may then be submitted as evidence of an accused persons guilt or innocence.

Where Google Earth Fits In

Google Earth (GE) can be used in multiple stages of the digital forensic process. Most often, you will find that it is used in the later parts of a case, when you need to analyze coordinates from various sources, or as a reporting tool to create presentations relating to geographic locations. In some cases, it may also be used to acquire GPS data from a device, although other tools may be more suited to collecting such data for a forensic investigation.

GPS Forensics

When a person uses a GPS device, he or she will enter in locations called *waypoints* that are stored in the GPS. The waypoint may be a person's current location, or a location that he or she wants to navigate to. The GPS device will use a series of waypoints to create a *route*, showing the person how to navigate from one location to others in a specific order. Because this information can be stored on the device, it can also be retrieved and examined during an investigation.

GPS devices will also store *tracks*, which are geographic points that the unit has been. When you turn on the GPS unit, it will connect to satellites and determine its current location. As you travel, additional track points will be stored as a record of where the GPS unit has been, and stored in a *track log*. By looking at the track log, you are able to view a listing of coordinates that the portable GPS has visited and, by extension, where its owner has been.

As we saw in Chapter 3, and revisit in the next chapter, Google Earth can be used to acquire data from a Garmin or Magellan GPS unit. In performing the import, you will see the number of waypoints, tracks and routes that are imported from a GPS device, which can then be reviewed in the 3D viewer.

However, importing GPS data in this way copies the data directly off of the device into Google Earth. It does not retrieve any data that may have been deleted, or is hidden on the device.

This can be a major issue if a particular location of interested a suspect visited existed in the deleted data,

and no longer appeared in the tracks you copied using Google Earth's import feature. For this reason, it is often best to use forensic tools to collect all of the data, not just what is visible to the device's interface, inclusive to any deleted or hidden data that may reside on the device.

Also, in acquiring the data from a GPS device for use with Google Earth, you want to ensure nothing is written to the GPS device. As the device will store files, your operating system or applications might write data without your knowledge or intention. If data from the original source of evidence has been modified, it could be challenged in court, and become inadmissible as evidence. To prevent this from happening, you should ensure that your forensic machine uses write protection and/or uses tools that are designed to gather evidence in a forensically sound manner, as we discuss in the next section.

TOOLS FOR RECOVERING EVIDENCE

As we have mentioned, it is important to recognize that GE is not a tool designed for digital forensic data collection. It will do a logical download of geolocation data, so anything that is been deleted from the device (i.e., waypoints, coordinates, etc.) will not be included when you use GE to import data from the device. To acquire data in a forensically sound manner, and get all the evidence that is available (regardless of whether it is deleted or hidden), more advanced tools should be considered.

In this section, we will discuss various tools that can be used to collect data from devices. There are software and hardware solutions that prevent your operating system or software like Google Earth from writing to the device or storage media, and ones that will create an exact duplicate so that you can work from an image of the data.

TIPS AND TRICKS

Working with Images and Other Copies of Data

By creating an image of what is stored on a computer or other devices, you are examining a copy of the data and not the original source. Forensic software that allows you to create an image in this way means that you can examine a computer or device without having to go through its operating system or user interface. In doing so, you are bypassing any passwords required to logon to a machine. Similarly, for mobile forensics, such tools can extract data while bypassing pattern locks, PINs or passwords.

Write Protection

Prior to acquiring data from a GPS unit with Google Earth, you should ensure that your forensic machine has USB write protection enabled. Because a GPS unit also can function as a mass storage device, it is essential to make sure that no data on the device is changed. Rather than simply plugging the GPS device into a USB port, you want to ensure that software write protection or a hardware write blocker is used to prevent any accidental modification of data.

Write blockers allow read commands to pass from a computer to a storage device, but block any write commands. In doing so, you can safely access the drive to view its contents and\or collect data. With a hardware blocker, the disk or device you are collecting evidence from plugs into a device that becomes a midway point between the forensic workstation and the storage you are acquiring data from. The ability to block writes may also be included in other forensic hardware tools that are used to image or duplicate the data on the suspect device.

There are also a number of software solutions that can be used to prevent your computer from writing to a storage device that you are collecting data from, such as a GPS device that is connected via a USB port. On a machine running Windows, you can use write protection software like:

- DSI USB Write Blocker (*document-solutions.biz/downloads/?did=9*)
- M2CFG USB Write Block (*www.m2cfg.com/usb_writeblock.htm*)
- NetWrix USB Blocker (*www.netwrix.com/usb_blocker_freeware.html*)
- Thumbscrew (*www.irongeek.com/i.php?page=security/thumbscrewsoftware-usb-write-blocker*)

There are also a number of tools for Mac computers that provide write protection, allowing you to safely acquire data, such as:

- Softblock (*www.blackbagtech.com/software-products/softblock-1/softblock.html*)
- Disk Arbitrator (*https://github.com/aburgh/Disk-Arbitrator/downloads*)

Tools Used to Acquire Evidence

In addition to the tools we have already mentioned, there are a number of products available for digital forensics investigations, which are commonly used by law enforcement and companies specializing in data collection. Using such suites of products, you will find that they have features and functions that will meet most of your needs throughout the process of acquiring, analyzing and reporting on digital evidence.

Guidance Software (*www.guidancesoftware.com*) is a company that creates a number of products used for digital forensics. The versions of EnCase are used to acquire evidence from hard drives, removable media (e.g., CDs, USB sticks, etc.), smartphones, tablets, GPS units and more. Using a GUI interface, the software can be used to acquire, analyze, and create reports to show what was found, where the data originated, details of files, and other pertinent facts that relate to your investigation. Once completed, you can have EnCase generate a report that can be provided to other investigators and the courts.

Cellebrite (*www.cellebrite.com*) is another company that is well known for its commercial digital forensic products. Using their software and hardware, you can acquire and examine data from mobile phones, GPS units, tablets, and other devices, as well as memory cards. The tools available can be used for manual acquisition, where there is a need to take screenshots or images of data, and for acquiring existing and deleted data from a device being examined.

Cellebrite also has tools specifically designed for investigations requiring the acquisition of data from GPS devices. Using these tools, you can extract data from portable GPS units like Tom Tom, Garmin and Mia, inclusive to any GPS fixes that may have been previously deleted. Once you have acquired the files using tools like Cellebrite and EnCase, you can then import them into Google Earth for further analysis.

File Converters

While you can import GPS data into Google Earth, you are limited to files for Garmin and Magellan units. If files have been retrieved from other types of GPS devices, then you will need to convert them prior to importing them into GE. Once converted to a Garmin or Magellan format or a KML file, you can then import the data into GE. Some of the file converters available include:

- GPSBabel (*www.gpsbabel.org*) is freeware application that runs on your computer, which converts waypoints, tracks and routes to different formats.
- GPS Visualizer (*www.gpsvisualizer.com/gpsbabel/*), which is a site that provides an online version of GPSBabel, allowing you to upload and convert the file on their site.
- TraceGPS (*www.tracegps.com/en/convert.htm*) is another site that allows you to upload and convert files from one format to another
- GPS Data Team (*http://tomtom.gps-data-team.com/poi/ov2-to-kml. php*), which is a site that can convert OV2 files used by Tom Tom GPS devices to a format used by Garmin devices.

DO YOU REALLY WANT TO DO THIS?

Just because you need the evidence does not mean that you should be the one to acquire it. Law enforcement may have a fulltime digital or computer forensic examiner, while a corporation or other organization may have someone on staff (such as in the I.T. department) who is trained in the collection of data using forensic methods and resources. Rather than doing the work yourself, you could have such a person collect the data for you, so you can work from a copy or image.

If you are not part of a formal investigation, you should ask *why* you are doing the work and where it might lead. Anyone using Google Earth has the ability to import and examine GPS data from a portable device, and retrieving and reviewing this information might be used for personal or non-investigative reasons. However, depending on what you find, that data may eventually become evidence in a court case, and how it was collected might be held to a higher standard. For example:

- A manager could import GPS data into Google Earth to review where an employee traveled during work hours. Is he or she traveling to meetings locations, customer offices and other work-related places, or visiting a bar or the beach? Looking at the GPS data would reveal where that employee goes, and if it was found the person was not doing their job, it could result in termination of employment. However, if the former employee challenged being fired and sued, then the data and methods of acquiring the GPS data could be questioned in civil litigation.
- If a friend was concerned that his/her spouse or significant other was cheating, you could examine where a portable GPS unit was taken in Google Earth. In doing so, you might confirm your friend's suspicions, but what if your findings became the basis for a divorce? What was a simple perusal of a person's goings on has now become evidence in a divorce case.

As you can see from these scenarios, a simple looksee can quickly change. When you acquire and examine any data, you should always assume that it could eventually become part of a criminal or civil case. Because of this, you should always try to follow best practices of data collection, documentation, and follow any procedures or policies created by your organization. By treating the acquisition of any data as a formal investigation, you will maintain good habits in the collection and analysis of evidence, and be prepared if you have to testify about it later.

ORGANIZING YOUR CASE

It is a good idea to make sure that when working on a geoforensic case in Google Earth, you make sure you keep your work organized so that it is easy to retrieve and share, that you can recover from mistakes and most importantly you can maintain consistent work flow. A recommended way to do this is to create case folders in Google Earth. It is suggested that an investigator create two types of folders when working a case in Google Earth:

- A case folder in the "My Places" top level directory for eventual case dissemination
- A "temporary" folder in the "Temporary Places" top level directory for experimenting and developing your work.

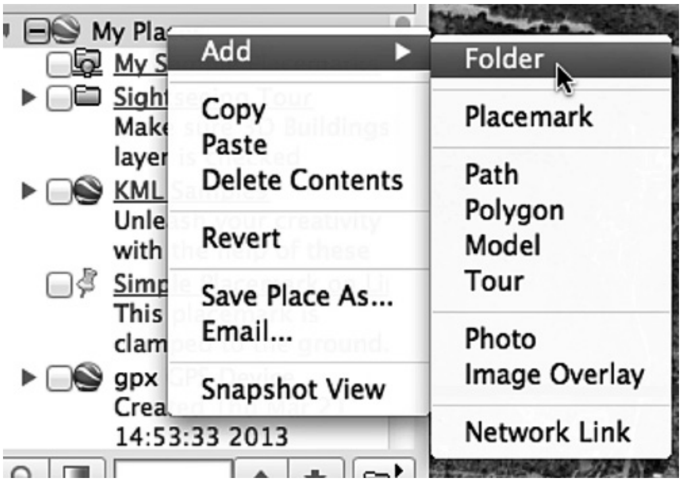

Figure 1. *Adding a folder*

Creating a folder in Google Earth is done by right-clicking on one of the top level directories, and when the context menu shown in Figure 1 appears, select *Add* and then click *Folder*.

Once you have created a folder, you are greeted with a dialog window to edit the settings of the folder. These settings are as follows.

- Name. Here is where you set the name of the folder. It is recommended that you use a consistent nomenclature for your particular organization. For instance <case name> – <case number>
- *Description*. You can give the folder a description of what is contains and a preview of this will appear below the folder. The description can also include links, photos and other HTML tags. This is covered in the previous chapters, and as well as Chapter 6.
- *Style, color*. This option becomes available once there are icons within the folder you are creating or any subfolders of the created folder. The option is used to create a universal color and label style in this folder and all its children.
- *View*. This option is used for creating one viewing angle for each of the placemarks contained in the folder. Once a view is set for a folder, double clicking on it will reset the view to match what was set. Setting the view will be covered in a section in Chapter 6.

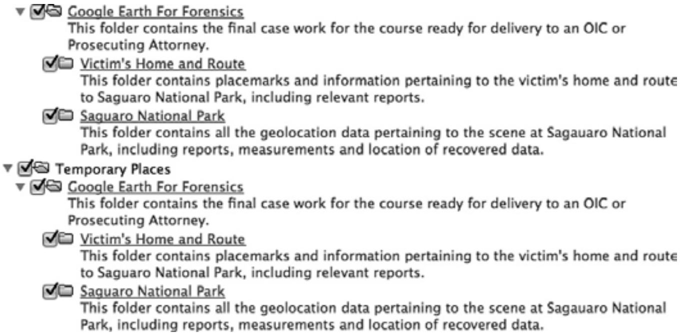

▾ ☑ Google Earth For Forensics
This folder contains the final case work for the course ready for delivery to an OIC or Prosecuting Attorney.
☑ Victim's Home and Route
This folder contains placemarks and information pertaining to the victim's home and route to Saguaro National Park, including relevant reports.
☑ Saguaro National Park
This folder contains all the geolocation data pertaining to the scene at Sagauaro National Park, including reports, measurements and location of recovered data.
▾ ☑ Temporary Places
▾ ☑ Google Earth For Forensics
This folder contains the final case work for the course ready for delivery to an OIC or Prosecuting Attorney.
☑ Victim's Home and Route
This folder contains placemarks and information pertaining to the victim's home and route to Saguaro National Park, including relevant reports.
☑ Saguaro National Park
This folder contains all the geolocation data pertaining to the scene at Sagauaro National Park, including reports, measurements and location of recovered data.

Figure 2. *Folder structure template*

In Chapter 6, we will work with a scenario to use the knowledge you have acquired throughout this book. For the purposes of our scenario for this course and to get you familiar with organizing your work, create the following structure by adding folders in *My Places* and *Temporary Places*. In using this template structure, it is encouraged that you change the template and narrative contained in the description to suit the needs of your agency (Figure 2).

Custom Icons

As we mentioned in Chapter 2, when creating placemarks, the *Style, Color tab* of the *Properties* dialog can be used to select a unique icon for each placemark. Using different icons makes your placemarks stand out from one another in the 3D viewer, and can provide an effective graphic representation of why a location is important and/or what was found there (e.g., a crime scene, remains, evidence, etc.).

As we will discuss in Chapter 6, you can select an icon from a library of icons that is included with Google Earth, or add a custom icon. Because you may find the ones included with GE limited, it may be useful to look at online resources, and take the time to choose ones that suit your purpose. A good site for custom icons is the Map Icons Collection (*http://mapicons.nicolasmollet.com*), which has hundreds of free icons that can be downloaded and used in your project. Other useful sites include:

- The Google Developers site (*http://code.google.com/p/googlemaps-icons/downloads/list*)
- Mapito (*http://www.mapito.net/map-marker-icons.html*)
- Benjamin Keen (*http://www.benjaminkeen.com/?p=105*)

Enabling Access to Local Files

Google Earth is set up natively to access the Internet to pull down content like map data or external files and pictures. But in digital forensics allowing Internet access by a program containing case data is generally considered to be a poor idea. It is of use, however, to use the capability of Google Earth to link to other files such as report PDFs or scene photographs. Below is the procedure for allowing Google Earth to link to files local to the examiner's machine (Figure 3).

1. From the *Tools* menu, and click the *Options* menu item (on a Mac click *Preferences*)
2. Click the *General* tab, and (as shown in the following figure) locate the *Placemark Balloons* section.
3. Click the *Allow access to local files and personal data* checkbox so it appears checked.
4. Accept the warning saying that access to local files might be risky, and click *OK*

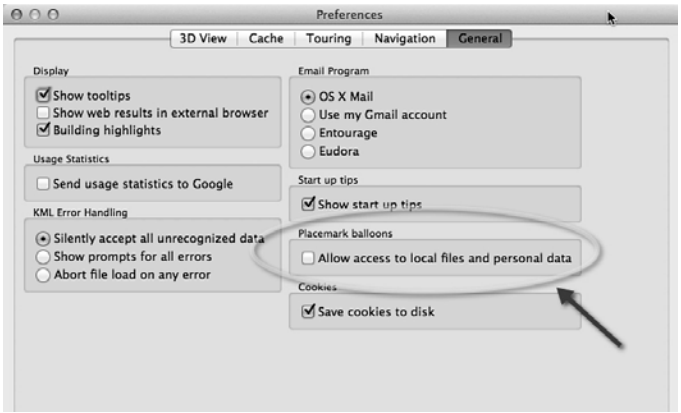

Display
☑ Show tooltips
□ Show web results in external browser
☑ Building highlights

Usage Statistics
□ Send usage statistics to Google

KML Error Handling
● Silently accept all unrecognized data
○ Show prompts for all errors
○ Abort file load on any error

Email Program
● OS X Mail
○ Use my Gmail account
○ Entourage
○ Eudora

Start up tips
☑ Show start up tips

Placemark balloons
□ Allow access to local files and personal data

Cookies
☑ Save cookies to disk

Figure 3. *Enable placemark balloon local access*

UNDERSTANDING WHAT YOU ARE LOOKING AT

When navigating through areas in Google Earth, it is important to realize that much of what is shown is not current. Some images may be recent, but others may be weeks, months or even years old. According to Google, most of the imagery you see is approximately 1–3 years old. As such, buildings that have been torn down may appear in GE, while those recently built are not visible. Similarly, the Street View does not contain real-time footage, so a familiar area may appear outdated as you take a virtual walk down the street. In using this tool, it is important to remember that what is displayed may not be an accurate representation of what is there now.

Why is He Blurry?

In Chapter 1, we mentioned that if you notice blurred imagery in GE, it may be due to slow or poor connections to the Internet. That being said, you can expect to see some blurred areas when viewing an area in Street View. To protect a person's privacy, Google uses an algorithm that will automatically blur a person's face and the license plates of vehicles so they cannot be identified.

Blocked Content

Generally, when you use Street View, you will not be able to access areas beyond the street. In other words, you will not be able to explore a mall's parking lot, private roads, empty fields, and so on. The reason for this is that Google uses a car with a panoramic camera on top of it to take photos as it drives down the street. It does not go off road to take photos, so you are limited to what is visible from the roadway. An exception to this is when a point of interest like Universal or Disney theme parks permit Google to enter and take digital photos of what is inside. Doing so allows you to take a virtual journey through that location.

Another time when you will notice missing content is when Google removes something that is considered inappropriate. An example of this is when you try and visit 105 Temperance Street in Manchester England, where you will find that you are prevented from navigating down a section of that roadway. The reason is that when the Google car drove by, the 15 lens panoramic camera captured multiple angles of a man and woman engaged in a sex act. The area was known for prostitution, and once it was discovered a salacious transaction had been photographed, Google blurred and later deleted the images.

Misinterpreted Content

While Google has captured unsavory and illegal acts on camera, and even used aerial imagery showing a crime scene, there are also times where people have mistakenly interpreted what is shown. An example of this occurred on Middle Road in St John's, Worcester, England when the Google car photographed a young girl lying face down in the road, with one shoe cast off in the gutter. When the images became available the next year, users of Google Maps and Google Earth were shocked to see what appeared to be a dead girl. Fortunately, things were not what they seemed. The 9-year-old was simply playing a prank on her friend, and had been unaware that Google had snapped her picture. Before you try looking for the imagery on Google Earth, you should be aware that they have already blurred and deleted images, preventing you from navigating down that road.

Removing Content

Problems related to what appears in Google Earth and Google Maps can be reported to the company, which may result in images being blurred, replaced or removed. To report an issue, you can use Google Maps (*https://maps.google.com*) to navigate to a particular location. Enter an address, and zoom into a location. When you are viewing a map or satellite image and spot a problem, you can click on the *Report a problem* link to display a dialog box that allows you to notify Google about incorrect road information, addresses, places, directions, or other issues. By clicking on the *Other Problems* link, you can report issues with satellite imagery, Street View, or other problems.

For Street View, anyone can report inappropriate content, or request that a location or person is blurred. Accessing Street View in Google Maps is the same as in Google Earth. You would navigate to a location and either zoom in as far as you can until it switches to Street View, or drop the pegman icon onto a location. Once you are in Street View, you will notice a you will see a *Report a problem* link in the lower right-hand corner. Upon clicking this, a separate browser window will open, where you can report inappropriate content. Once this window opens, you will see a picture of what you were looking at in Street View, which you can adjust to focus on a particular part of the image. You can then request that a face, your home, car or license plate, or a different object is blurred. While you have reported the issues using Google Maps, the changes will also appear in Google Earth.

ABOUT THE AUTHOR

Michael Cross (MCSE, MCP+I, CNA, Network+) is an Internet specialist/ computer forensic analyst with the Niagara Regional Police Service (NRPS). He performs computer forensic examinations on computers involved in criminal investigation. He also has consulted and assisted in cases dealing with computer-related/Internet crimes. In addition to designing and maintaining the NRPS Web site at www.nrps.com and the NRPS intranet, he has provided support in the areas of programming, hardware, and network administration. As part of an information technology team that provides support to a user base of more than 800 civilian and uniform users, he has a theory that when the users carry guns, you tend to be more motivated in solving their problems. Michael also owns KnightWare (www.knightware.ca) that provides computerrelated services such as Web page design, and Bookworms (www.bookworms.ca), where you can purchase collectibles and other interesting items online. He has been a freelance writer for several years, and he has been published more than three dozen times in numerous books and anthologies. He currently resides in St. Catharines, Ontario, Canada, with his lovely wife, Jennifer, his darling daughter, Sara, and charming son, Jason.

ABOUT THE AUTHOR

Michael Harrington is a former law enforcement officer with over 10 years of experience in digital forensics. He lectures on mobile forensics around the world and has been involved in various forensic projects including Pandora's Box and WOLF. Michael has been published in the Thomas J Cooley Law Journal and on Forensic Focus. He also writes on the subject of mobile forensics at http://mobileforensics.wordpress.com/.

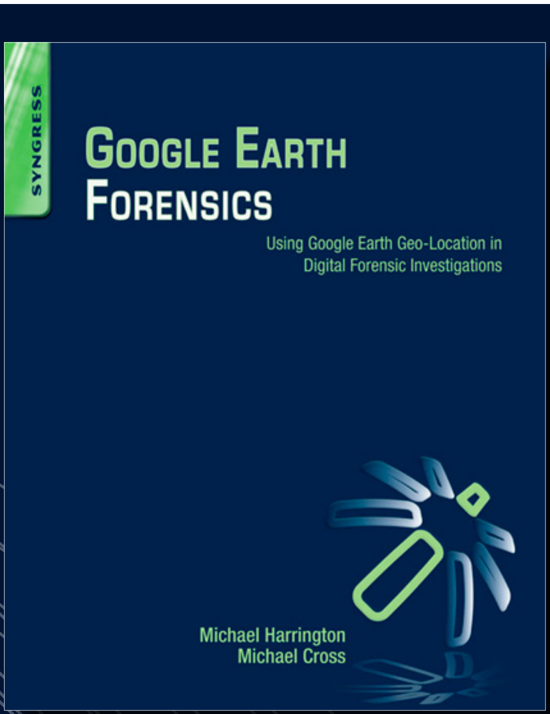
SYNGRESS
GOOGLE EARTH
FORENSICS
Using Google Earth Geo-Location in Digital Forensic Investigations
Michael Harrington
Michael Cross

Could Turn the Engines off at 35,000 Feet

ROB SOMERVILLE

In my previous column, I highlighted the hidden threats that technology can engender if the supply chain is compromised. Sadly, with the detention and interrogation of Chris Roberts of the Colorado-based One World Labs, even exposing the more obvious threats seems to bring paranoia and panic as a response.

So here we go again – A security expert who has identified weaknesses in the in-flight entertainment systems of aircraft that to quote Roberts "could turn the engines off at 35,000 feet" has been pulled off a flight, no doubt detained in some dreary room at Syracuse airport to be grilled about his knowledge. To add insult to injury, the FBI forensically examined the plane to check if anything had been compromised. All ironically in the face of Roberts' TSA (Transportation Security Administration) clearance and pro-actively working with the intelligence community to expose and help mitigate these type of risks. Maybe it was the recent article on Fox News that brought this disproportional response[1], but it is clearly apparent that the established order are acting like an injured animal biting the hand of the veterinarian trying to treat its wounds. Welcome to the particularly non-exclusive club of the persecuted whistle-blower Chris, as you have discovered to your cost the idiom of "No good deed goes unpunished" is alive and well in the 21st century.

Where the decision to take this action originated from is unclear, so it would be unfair to blame any of the alphabet and law enforcement agencies involved. No doubt if pressed, the Nuremberg defence will be rolled out again – after all, orders are orders and the responsibility lies further up the chain of command. It is a pity that 70 years after the close of the Second world war we still struggle with the unresolved tensions of personal versus corporate and institutional responsibility. Ironically, the larger and more expansive the organisation, the more dilute freedom to make autonomous decisions becomes and to take any responsibility harder still. After all, it is the first duty of the professional to toe the line when it comes to policy, and any dissent is looked upon as insubordination, disloyalty or just sheer rebellion. And we wonder why there are so many injustices and inequalities – faced with the institutionalised tar pit of policy, box ticking exercises, "lessons being learned" and enquiries that are loaded from the start, it is little wonder that few brave souls muster sufficient courage and gumption to raise their heads above the parapet. And of course, there is always the added benefit of being branded a crank, having your motives criticised or even more worryingly a concerted attack on your character and exposure of past mistakes by the media or the intelligence services. Dig deep enough and you will find dirt on anyone, but it takes a particular form of pond-life to twist this reality and use it to obscure the fundamental truth that the credible whistle-blower is attempting to expose. Hopefully, Chris Roberts will be rewarded with an apology, but I somewhat suspect that if this story gains more traction the tactic of shooting the messenger will be brought into play.

Unequivocally, airline safety must be a top priority for governments and the industry to address. Flying at 35,000 feet in an aluminium tube constructed by the lowest bidder (Thanks Steve Buscemi of Armageddon fame for the quote) has an unnerving effect on many. The industry is notoriously sensitive to wars, terrorist action, economic conditions and global geo-politics negatively affecting revenue and profit streams. The profit margins per passenger are generally extremely low, so anything that dents consumer confidence can have a major effect on the bottom line. The trend for corporates implementing home working and telecommuting has created a schism within the industry where major investment has been made to attract first-class travellers (with their corresponding large wallets) and a reduction or elimination of business class accommodation, leading to an almost 2 tier system. The remaining bulk of passengers are not frequent flyers, and the substantial proportion of this travel is closely knit to the vacation and travel industry. Any panic can potentially lead an airline to bankruptcy as has happened many times in the past, with the corresponding financial carnage filtering down to hotel chains, car rental businesses and small family-run hotels. Despite the irony that it is statistically more dangerous to get in a car or cross the road than fly across the world, perceptions are fickle so it is understandable from the commercial angle (although unforgivable from a risk management or ethical perspective) as to why such a hard line is being taken.

Good risk management at its core must foster a culture of trust and openness, if not reward. While it is beyond argument that certain sensitive details should be kept well away from parties that could exploit these weaknesses for gain, the old problem of leverage raises its ugly head. What do you do if a vulnerability is discovered and you inform the software developer or manufacturer confidentially but they bury their heads in the sand? Do you go public? Quietly reassure yourself that you have done your best but let it lie? Go to the top of the organisation in the hope that those in a position of responsibility will act? Large organisations suffer from a peculiar type of inertia when facing such crises, and the results are not pretty – especially when you have a business sector so closely regulated by government. Batten down the hatches, wear Teflon coating and pray that the problem goes away.

This is not the way it should be. Either the airline industry wants the help of the white hats and as a result takes them and their analysis seriously, or undermines genuine efforts to be a part of the solution and becomes a part of the problem itself – immediately playing into the hands of the black hats by giving them strategic advantage. Never so true is the maxim "It takes one to know one" more relevant. All the policies, governance and PR machinery available will not identify security weaknesses, only the sharp-eyed forensics and security analyst with experience, the relevant technology and *carte blanche* to examine the organisation as a whole from every conceivable angle. This incident is more reminiscent of a "Blue on Blue" or "Friendly fire" scenario than one that reinforces trust and partnership between business, white hats and law enforcement repelling a common enemy. This is the danger where an offence is classified as strict liability – the law does not consider the *mens rea* or state of mind of the defendant. Those concerned may think that the removal of Chris Roberts from the flight was judicious and proportionate, but others may have second thoughts when it comes to reporting issues from now on. Under similar circumstances, rather than approaching the media, maybe the better approach is to befriend a long-time experienced airline pilot – like few others they understand the true risks of flying. Or if you are of a more delicate composition, just quietly forget about it.

I have no doubt that One World Labs and Chris Roberts in particular will take this breach of common sense on the chin. It goes with the territory, for the world of intelligence services, law enforcement, forensic and penetration security traditionally attract knee jerk and emotional reactions from the uninformed. But until the commercial sector has the courage to embrace bad news exposed at even the most public systems level, it will continue to find that solutions to problems buried much deeper below the surface will evade their grasp. Despite years of denial even the behemoth Microsoft has realised that this is not a sustainable long term strategy.

References

- *http://www.foxnews.com/us/2015/04/17/security-expert-pulled-off-flight-by-fbi-after-exposing-airline-tech*

ABOUT THE AUTHOR

Rob Somerville has been passionate about technology since his early teens. A keen advocate of open systems since the mid-eighties, he has worked in many corporate sectors including finance, automotive, airlines, government and media in a variety of roles from technical support, system administrator, developer, systems integrator and IT manager. He has moved on from CP/M and nixie tubes but keeps a soldering iron handy just in case.

Attend

InterDrone

The International Drone Conference and Exposition

InterDrone is Three Awesome Conferences:

Drone TECHCON
For Builders

More than 35 classes, tutorials and panels for hardware and embedded engineers, designers and software developers building commercial drones and the software that controls them.

Drone FLYER
For Flyers and Buyers

More than 35 tutorials and classes on drone operations, flying tips and tricks, range, navigation, payloads, stability, avoiding crashes, power, environmental considerations, which drone is for you, and more!

Drone BUSINESS
For Business Owners, Entrepreneurs & Dealers

Classes will focus on running a drone business, the latest FAA requirements and restrictions, supporting and educating drone buyers, marketing drone services, and where the next hot opportunities are likely to be!

The Largest Commercial Drone Show in North America

Meet with 80+ exhibitors!
Demos! Panels! Keynotes!
The Zipline!

September 9-10-11, 2015
Rio, Las Vegas
www.InterDrone.com

A BZ Media Event