identv™

SHAG I

11/10/2017

Installation and Configuration for Inference Engine

Ubuntu / OS Mirrors

Memory Mapped File

Technical Documentation Team
IDENTV, LLC

SHAGI

Table of Contents

1 Copyright, Trademarks, Notice of Proprietary Rights

11 COPYRIGHT

Information in this document may change without notice and does not represent a commitment on the part of IDenTV, LLC.

IDenTV, LLC. claims copyright in this documentation as an unpublished work, revisions of which were first licensed on the date indicated in the foregoing notice. The claim of copyright does not imply waiver of IDenTV's other rights. See notice of Proprietary Rights.

12 TRADEMARKS

The IDenTV logos are registered trademarks of IDenTV, LLC.

All other company and product names used herein may be the trademarks or registered trademarks of their respective companies.

13 NOTICE OF PROPRIETARY RIGHTS

This documentation is a confidential trade secret and the property of IDenTV, LLC. and contains IDenTV's Confidential, Proprietary, or Privileged Information. Use, examination, reproduction, copy, de-compilation, transfer, and/or disclosure to others are strictly prohibited except by express written agreement with IDenTV, LLC.

2. Revision History

Date	Version	Author	Document Changes
11/10/2017	1.0	Stephen	v.1.0
11/17/2017	1.0	Stephen	Change Title to include Ubuntu / OS Mirrors

Table 1: Revision History

3. Approvals

Role	Name	Title	Signature	Date
Project Sponsor:	ECS			
Project Manager:	Nikki Maquindang	Project Manager		
System Architect:	Erick Wipprecht			
System Architect:	Marco Baena			
System Architect:	Mark Hughes			
Development Lead:				
User Experience				
Quality Lead:				
Content Lead:	Stephen Mitchell	Sr. Technical	SWM	11/10/2017

Table 2: Approvals

4. Points of Contact

Name	Title	Email	Contact Number
Erick Wipprecht		erick@identv.com	
Marco Baena		marco@identv.com	
Mark Hughes		mark@identv.com	
Nikki Maquindang	Project Manager	nikki@identv.com	702.460.5622

Table 3: Points of Contact

4.1 Help Desk

In the event of a question about the operation of the systems, users should contact the office of Nikki Maquindang at (702) 460-5622, or via email nikki@identv.com

5. Installation and Configuration of the Inference Engine

5.1 Libraries and Requirements:

It is necessary to install and configure the following libraries.

5.11 Required Libraries

1. Python 2.7 – should be installed by default on Ubuntu
2. NVidia Graphics Driver 3.84 or higher
3. Cuda 8.0 (tested on 8.0 although 9.0 should be fine also):
   https://developer.nvidia.com/cuda-toolkit
4. CUDNN 6.0 (tested on 6.0 but 7.0 should be fine with Cuda 9.0):
   https://developer.nvidia.com/cudnn
5. OpenCV (latest version)
   5.1. Apt-Get: sudo apt-get install libopencv-dev python-opencv
   5.2. Build from scratch: See: Figure 1
6. Tensorflow GPU: sudo pip install tensorflow-gpu
7. Keras: pip install --upgrade keras
8. Pytorch:
   8.1. http://download.pytorch.org/whl/cu80/torch-0.2.0.post3-cp27-cp27mu-manylinux1_x86_64.whl
   8.2. pip install torchvision pip install
9. PIL: sudo pip install Pillow
10. Lxml: sudo pip install lxml

6. Inference Engine

The Inference Engine itself is written in python and can be run from the command line. To integrate with the ECS system, the Inference Engine consists of two main components:

1. The proxy file buffering service to integrate with ECS' platform on TIPS
2. The python inference engine based on Tensorflow

Note

This engine ingests an image one at a time from a local file system which is generated by the proxy system that connects to ECS' platform. Each file is analyzed by the inference engine, and the output is inserted into the same file system accessed by the Proxy service. The output format is straightforward and consists of a list of detections as generated by the inference engine. These detections include:

- Classname: Class of the object detected
- Confidence: How confident the model is with the detection
- X1: The leftmost position of the bounding box
- Y1: The topmost position of the bounding box
- X2: The rightmost position of the bounding box
- Y2: The bottommost position of the bounding box

Many detections can be generated from every single image; therefore, the output consists of a list format.

To run the Inference Engine, 'cd' to the **Inference Engine** folder and type '**nohup python InferenceEngine.py &**'. This will initialize and start the Inference Engine which will start polling the local file system for images coming from the proxy component.

```
wget https://github.com/opencv/opencv/archive/2.4.12.zip
sudo apt-get update
sudo apt-get install unzip
unzip 2.4.12.zip
cd opencv-2.4.12/
mkdir build
sudo apt-get install build-essential python-dev gfortran
sudo apt-get install python-numpy python-scipy python-matplotlib ipython
ipython-notebook python-pandas python-sympy python-nose
sudo apt-get install build-essential cmake pkg-config
sudo apt-get install libjpeg8-dev libtiff5-dev libjasper-dev libpng12-dev
sudo apt-get install libavcodec-dev libavformat-dev libswscale-dev libv4l-dev
sudo apt-get install libxvidcore-dev libx264-dev
sudo apt-get install libgtk-3-dev
sudo apt-get install libatlas-base-dev gfortran
sudo apt-get install python2.7-dev
cd build/
cmake -D CMAKE_BUILD_TYPE=RELEASE -D
CMAKE_INSTALL_PREFIX=/usr/local -D INSTALL_PYTHON_EXAMPLES=ON -D
BUILD_opencv_gpu=OFF -D INSTALL_C_EXAMPLES=OFF -D BUILD_EXAMPLES=ON..
make -j4
sudo make install
sudo ldconfig
```

Figure 1: Commands to build opencv 2.4.13 from scratch on Ubuntu 16.04

7. Memory Mapped File

7.1 Integration details

This section describes the format for exchanging data between the Inference Engine and the software calling it.

7.1.1 Format definition:

- The header size is 32 bytes.
- The image data starts right after the header.

7.1.2 Constants used when writing the headers:

- name | value | type (size) STATE_WITH_DATA | 2 | byte (1)
- STATE_READY_FOR_WRITING | 1 | byte (1)
- ENDIAN_TESTING_VALUE | 555 | short (2)

7.1.3 Header format:

- 0 | sync flag | byte (1) | This value is used to signal the state of the file: STATE_WITH_DATA / STATE_READY_FOR_WRITING.
- 1 | ENDIAN_TESTING_VALUE | short (2) | Used to validate that the other end is writing with the same endianness.
- 3 | imageDataSize | integer (4) | The size in bytes of the image data. 7 | columns | short (2) | The width of the image in pixels.
- 9 | rows | short (2) | The height of the image in pixels.
- 11 | channels | short (2) | The number of channels in the image. This field is not used. (The image must have three channels.)
- 13 | depth | short (2) | The color depth. Not used.
- 15 | type | short (2) | The image type. This field is not used.

Note

Position 0 holds the first byte.

8. Image Data

The pixels are written directly after the header. The size in the header states the total length. The length value can be validated by computing: rows x columns x 3 = image size

9. Output Data

When the Inference Engine (IE) writes the output, it first stores the number of detects in position 3.

Then it writes the detect strings consecutively right after that number. The detect strings have a fixed size of 170 bytes.

Position | **Name** | **Type** (**size**) | **Description**

- 3 | number-of-detects | short (2) | Number of detects. This will be the number of detection strings to be read.
- 5 | detects-block | bytes (number-of-detects x 170)

The detect-string format is: [<className>;<certainty>;<x1>;<y1>;<x2>;<y1>]

Where all the values are written in readable (ASCII) text, being:

className: is the name defining the class type detected (e.g., "CAR"). It is recommended that the className value is not longer than 64 bytes.

Certainty: is the percentage of certainty for that detect (e.g., 0.95). It is recommended that the certainty value is not longer than 5 bytes.

- x1: The upper left X coordinate for the bounding box.
- y1 The upper left Y coordinate for the bounding box.
- x2 The lower right X coordinate for the bounding box.
- y2 The lower right Y coordinate for the bounding box.

Notes

- The detect string is supposed to be padded with bytes holding the value 0 (actual 0, not the ASCII value for char '0').
- At least the first byte (in this sequence holding zeroes (0) is mandatory to maintain the integration agnostic to the programming language and operating system. This is the equivalent to ending a string with a 0 in C++.
- The string with the values (before the zeroes) must be enclosed in brackets: "[" and "]"

Example: [CAR;0.98;100;50;300;150] (here go the zeros)

10. Notes on the Flow

The memory mapped file should be set to **STATE_READY_FOR_WRITING** (**1**). The inference engine should be started before the calling process.

10.1 The Flow:

Assuming the memory mapped file is set to **STATE_READY_FOR_WRITING** (**1**)**.**

1. Start the Inference Engine. It will load resources and hold until the state changes (first byte in the mapped file).
2. Start the injection software (e.g., TIPS).

The Injection Software (IS) will read the state and trigger the writing of the image (it also reads the empty output that it disregards).

- The IS will write **STATE_WITH_DATA** and hold until it changes.
- The Inference Engine (IE) will see the change in the state and read the image.
- The IE will process the image, write the possible output and write the state STATE_READY_FOR_WRITING.
- The IS will read the output and start the process over.

11. Ubuntu / OS Mirrors:

This is a dump of the repos added to the GPU nodes for the initial training engine installation:

root@chrome-i-gpu-node-5:/etc/apt/sources.list.d# for file in `ls *.list`; do echo $file; cat $file; echo;done
cuda.list
deb http://developer.download.nvidia.com/compute/cuda/repos/ubuntu1604/x86_64 /

graphics-drivers-ubuntu-ppa-xenial.list

deb http://ppa.launchpad.net/graphics-drivers/ppa/ubuntu xenial main
deb-src http://ppa.launchpad.net/graphics-drivers/ppa/ubuntu xenial main

identv.list

deb https://repo.identv.com identv main

> **Note**
> (I'm guessing we won't want to mirror this, but rather provide and install the small number of packages manually)

saltstack.list

deb https://repo.saltstack.com/apt/ubuntu/16.04/amd64/latest xenial main

(including this more for info – with the small number of servers in play salt isn't an asset)

webupd8team-ubuntu-java-xenial.list

deb http://ppa.launchpad.net/webupd8team/java/ubuntu xenial main

deb-src http://ppa.launchpad.net/webupd8team/java/ubuntu xenial main

PyPI Mirror

We'll need to mirror python packages for pip to work properly. These links should be helpful:

https://www.pypi-mirrors.org/

https://pypi.python.org/pypi/bandersnatch

Node.js Mirror

We'll also need to mirror packages installed via npm. I believe this is only to support the tagger. I'm honestly not sure how to best approach this as this isn't something I've done previously and am not sure how mirrors are structured (one big mirror that **npm** pulls from? Packages potentially residing on separate mirrors?)

Currently installed **npm** packages are listed below, both in user and global contexts. I'm not sure if there's overlap or not – the global list may be complete.

$ cat global.txt

/usr/lib

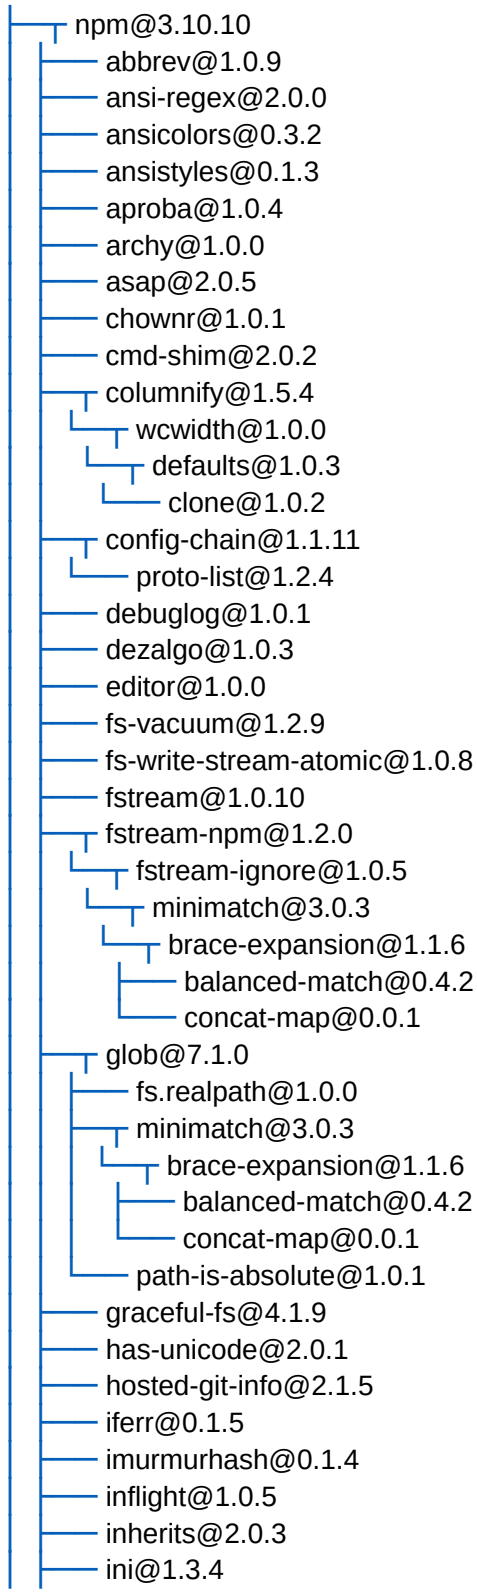

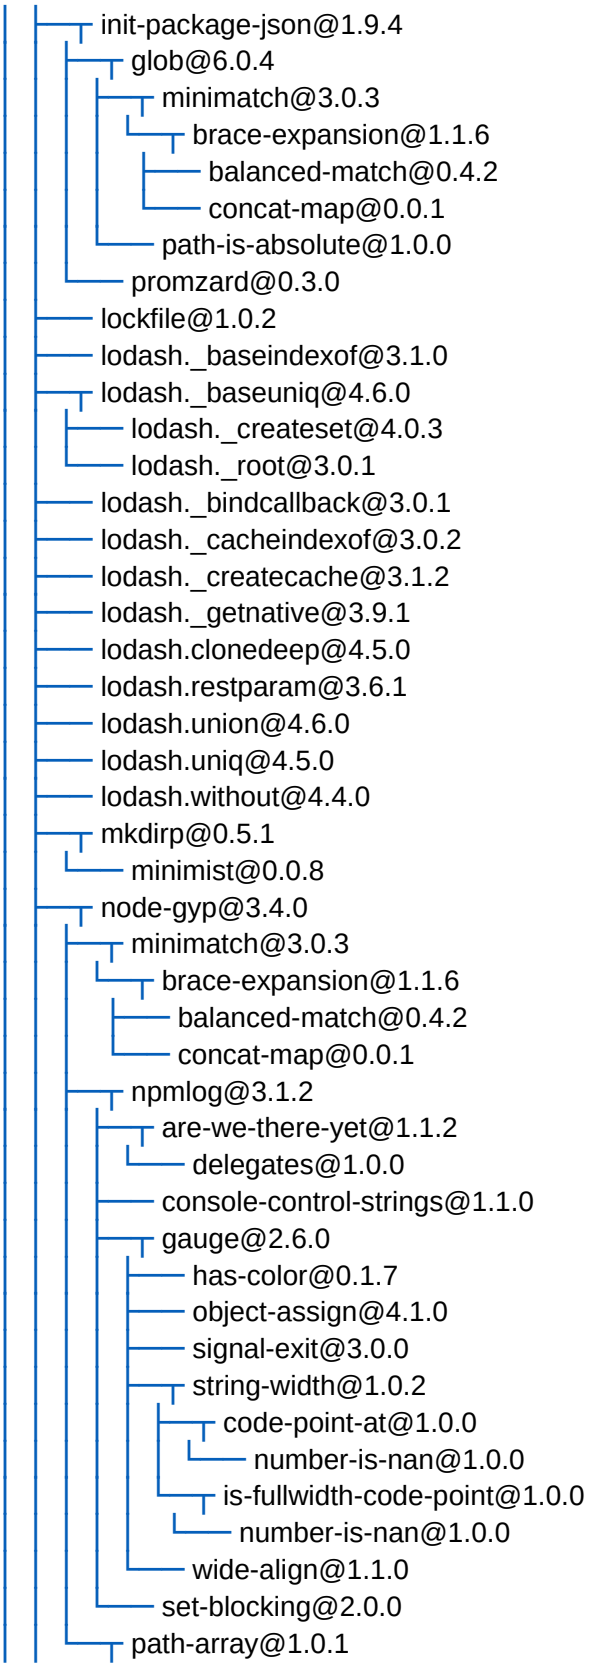
init-package-json@1.9.4
glob@6.0.4
minimatch@3.0.3
brace-expansion@1.1.6
balanced-match@0.4.2
concat-map@0.0.1
path-is-absolute@1.0.0
promzard@0.3.0
lockfile@1.0.2
lodash._baseindexof@3.1.0
lodash._baseuniq@4.6.0
lodash._createset@4.0.3
lodash._root@3.0.1
lodash._bindcallback@3.0.1
lodash._cacheindexof@3.0.2
lodash._createcache@3.1.2
lodash._getnative@3.9.1
lodash.clonedeep@4.5.0
lodash.restparam@3.6.1
lodash.union@4.6.0
lodash.uniq@4.5.0
lodash.without@4.4.0
mkdirp@0.5.1
minimist@0.0.8
node-gyp@3.4.0
minimatch@3.0.3
brace-expansion@1.1.6
balanced-match@0.4.2
concat-map@0.0.1
npmlog@3.1.2
are-we-there-yet@1.1.2
delegates@1.0.0
console-control-strings@1.1.0
gauge@2.6.0
has-color@0.1.7
object-assign@4.1.0
signal-exit@3.0.0
string-width@1.0.2
code-point-at@1.0.0
number-is-nan@1.0.0
is-fullwidth-code-point@1.0.0
number-is-nan@1.0.0
wide-align@1.1.0
set-blocking@2.0.0
path-array@1.0.1

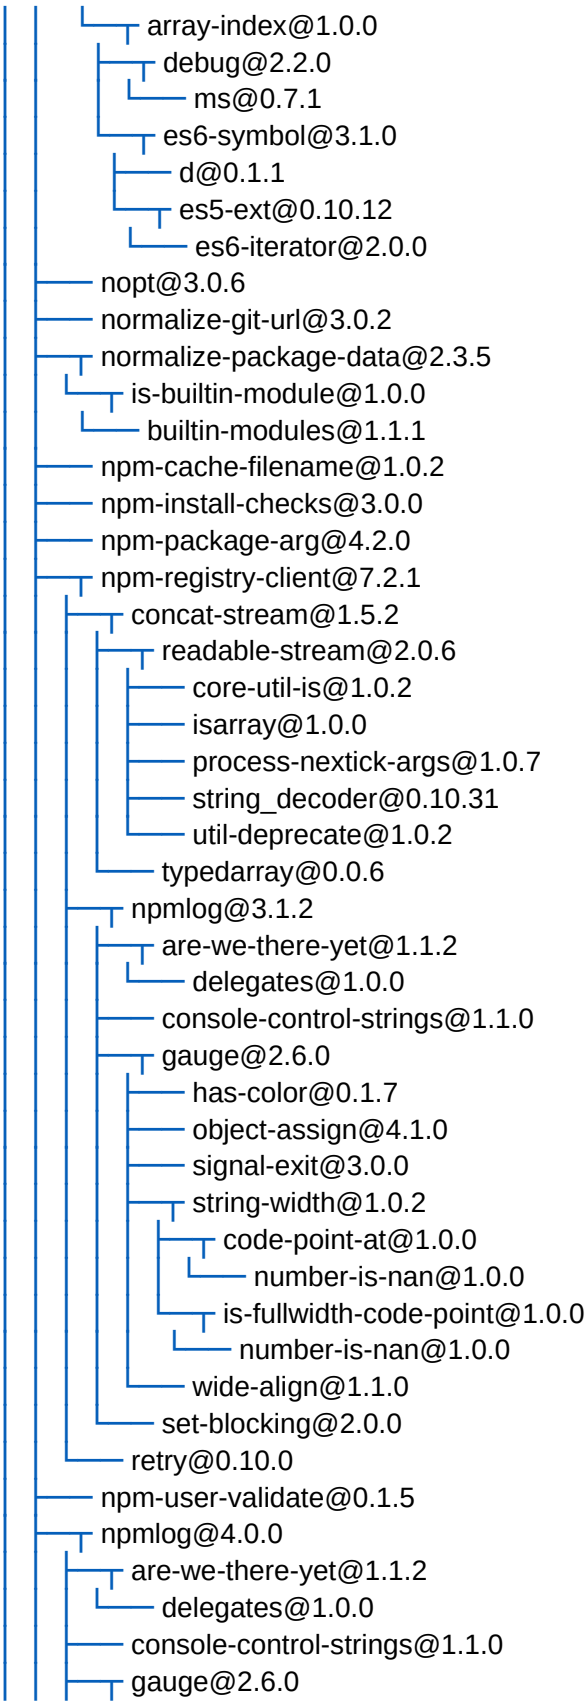
array-index@1.0.0
debug@2.2.0
ms@0.7.1
es6-symbol@3.1.0
d@0.1.1
es5-ext@0.10.12
es6-iterator@2.0.0
nopt@3.0.6
normalize-git-url@3.0.2
normalize-package-data@2.3.5
is-builtin-module@1.0.0
builtin-modules@1.1.1
npm-cache-filename@1.0.2
npm-install-checks@3.0.0
npm-package-arg@4.2.0
npm-registry-client@7.2.1
concat-stream@1.5.2
readable-stream@2.0.6
core-util-is@1.0.2
isarray@1.0.0
process-nextick-args@1.0.7
string_decoder@0.10.31
util-deprecate@1.0.2
typedarray@0.0.6
npmlog@3.1.2
are-we-there-yet@1.1.2
delegates@1.0.0
console-control-strings@1.1.0
gauge@2.6.0
has-color@0.1.7
object-assign@4.1.0
signal-exit@3.0.0
string-width@1.0.2
code-point-at@1.0.0
number-is-nan@1.0.0
is-fullwidth-code-point@1.0.0
number-is-nan@1.0.0
wide-align@1.1.0
set-blocking@2.0.0
retry@0.10.0
npm-user-validate@0.1.5
npmlog@4.0.0
are-we-there-yet@1.1.2
delegates@1.0.0
console-control-strings@1.1.0
gauge@2.6.0

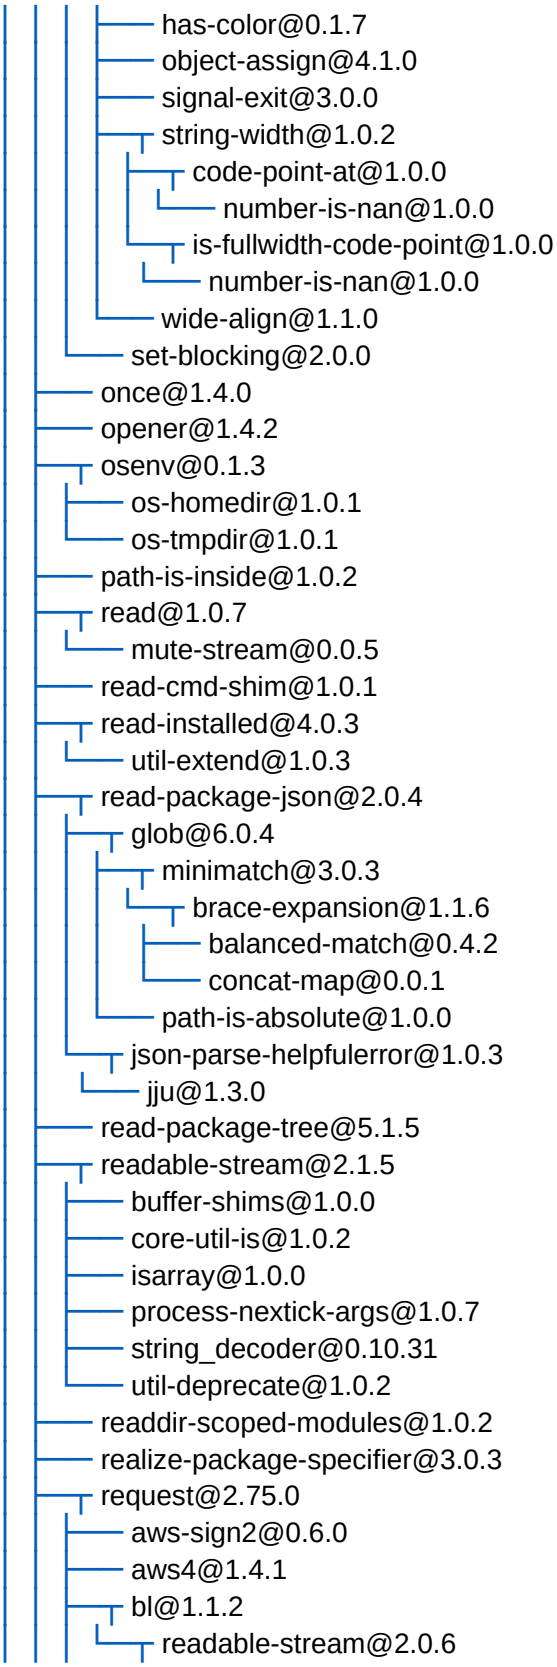
has-color@0.1.7
object-assign@4.1.0
signal-exit@3.0.0
string-width@1.0.2
code-point-at@1.0.0
number-is-nan@1.0.0
is-fullwidth-code-point@1.0.0
number-is-nan@1.0.0
wide-align@1.1.0
set-blocking@2.0.0
once@1.4.0
opener@1.4.2
osenv@0.1.3
os-homedir@1.0.1
os-tmpdir@1.0.1
path-is-inside@1.0.2
read@1.0.7
mute-stream@0.0.5
read-cmd-shim@1.0.1
read-installed@4.0.3
util-extend@1.0.3
read-package-json@2.0.4
glob@6.0.4
minimatch@3.0.3
brace-expansion@1.1.6
balanced-match@0.4.2
concat-map@0.0.1
path-is-absolute@1.0.0
json-parse-helpfulerror@1.0.3
jju@1.3.0
read-package-tree@5.1.5
readable-stream@2.1.5
buffer-shims@1.0.0
core-util-is@1.0.2
isarray@1.0.0
process-nextick-args@1.0.7
string_decoder@0.10.31
util-deprecate@1.0.2
readdir-scoped-modules@1.0.2
realize-package-specifier@3.0.3
request@2.75.0
aws-sign2@0.6.0
aws4@1.4.1
bl@1.1.2
readable-stream@2.0.6

```
        core-util-is@1.0.2
        isarray@1.0.0
        process-nextick-args@1.0.7
        string_decoder@0.10.31
        util-deprecate@1.0.2
    caseless@0.11.0
    combined-stream@1.0.5
      delayed-stream@1.0.0
    extend@3.0.0
    forever-agent@0.6.1
    form-data@2.0.0
      asynckit@0.4.0
    har-validator@2.0.6
      chalk@1.1.3
        ansi-styles@2.2.1
        escape-string-regexp@1.0.5
        has-ansi@2.0.0
        supports-color@2.0.0
      commander@2.9.0
        graceful-readlink@1.0.1
      is-my-json-valid@2.15.0
        generate-function@2.0.0
        generate-object-property@1.2.0
          is-property@1.0.2
        jsonpointer@4.0.0
        xtend@4.0.1
      pinkie-promise@2.0.1
        pinkie@2.0.4
    hawk@3.1.3
      boom@2.10.1
      cryptiles@2.0.5
      hoek@2.16.3
      sntp@1.0.9
    http-signature@1.1.1
      assert-plus@0.2.0
      jsprim@1.3.1
        extsprintf@1.0.2
        json-schema@0.2.3
        verror@1.3.6
      sshpk@1.10.1
        asn1@0.2.3
        assert-plus@1.0.0
        bcrypt-pbkdf@1.0.0
        dashdash@1.14.0
        ecc-jsbn@0.1.1
```

```
│ │ │   ├── getpass@0.1.6
│ │ │   ├── jodid25519@1.0.2
│ │ │   ├── jsbn@0.1.0
│ │ │   └── tweetnacl@0.14.3
│ │ ├── is-typedarray@1.0.0
│ │ ├── isstream@0.1.2
│ │ ├── json-stringify-safe@5.0.1
│ │ ├─┬ mime-types@2.1.12
│ │ │ └── mime-db@1.24.0
│ │ ├── node-uuid@1.4.7
│ │ ├── oauth-sign@0.8.2
│ │ ├── qs@6.2.1
│ │ ├── stringstream@0.0.5
│ │ ├── tough-cookie@2.3.1
│ │ └── tunnel-agent@0.4.3
│ ├── retry@0.10.0
│ ├── rimraf@2.5.4
│ ├── semver@5.3.0
│ ├── sha@2.0.1
│ ├── slide@1.1.6
│ ├── sorted-object@2.0.1
│ ├── strip-ansi@3.0.1
│ ├─┬ tar@2.2.1
│ │ └── block-stream@0.0.8
│ ├── text-table@0.2.0
│ ├── uid-number@0.0.6
│ ├── umask@1.1.0
│ ├─┬ unique-filename@1.1.0
│ │ └── unique-slug@2.0.0
│ ├── unpipe@1.0.0
│ ├─┬ validate-npm-package-license@3.0.1
│ │ ├─┬ spdx-correct@1.0.2
│ │ │ └── spdx-license-ids@1.2.0
│ │ └─┬ spdx-expression-parse@1.0.2
│ │   ├── spdx-exceptions@1.0.4
│ │   └── spdx-license-ids@1.2.0
│ ├─┬ validate-npm-package-name@2.2.2
│ │ └── builtins@0.0.7
│ ├─┬ which@1.2.11
│ │ └── isexe@1.1.2
│ ├── wrappy@1.0.2
│ └── write-file-atomic@1.2.0
└─┬ pm2@2.7.1
  ├─┬ async@2.5.0
  │ └── lodash@4.17.4
```

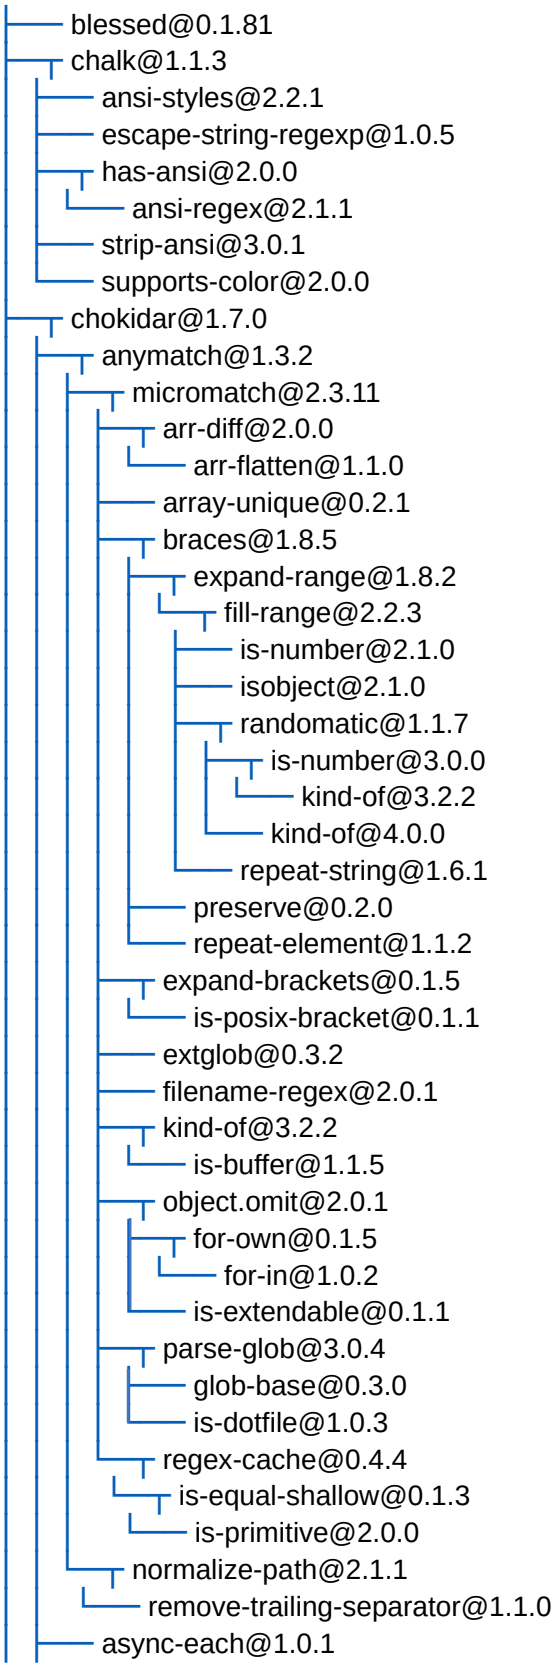
blessed@0.1.81
chalk@1.1.3
ansi-styles@2.2.1
escape-string-regexp@1.0.5
has-ansi@2.0.0
ansi-regex@2.1.1
strip-ansi@3.0.1
supports-color@2.0.0
chokidar@1.7.0
anymatch@1.3.2
micromatch@2.3.11
arr-diff@2.0.0
arr-flatten@1.1.0
array-unique@0.2.1
braces@1.8.5
expand-range@1.8.2
fill-range@2.2.3
is-number@2.1.0
isobject@2.1.0
randomatic@1.1.7
is-number@3.0.0
kind-of@3.2.2
kind-of@4.0.0
repeat-string@1.6.1
preserve@0.2.0
repeat-element@1.1.2
expand-brackets@0.1.5
is-posix-bracket@0.1.1
extglob@0.3.2
filename-regex@2.0.1
kind-of@3.2.2
is-buffer@1.1.5
object.omit@2.0.1
for-own@0.1.5
for-in@1.0.2
is-extendable@0.1.1
parse-glob@3.0.4
glob-base@0.3.0
is-dotfile@1.0.3
regex-cache@0.4.4
is-equal-shallow@0.1.3
is-primitive@2.0.0
normalize-path@2.1.1
remove-trailing-separator@1.1.0
async-each@1.0.1

```
├─┬ UNMET OPTIONAL DEPENDENCY fsevents@^1.0.0
├── glob-parent@2.0.0
├── inherits@2.0.3
├─┬ is-binary-path@1.0.1
│ └── binary-extensions@1.10.0
├─┬ is-glob@2.0.1
│ └── is-extglob@1.0.0
├── path-is-absolute@1.0.1
└─┬ readdirp@2.1.0
  ├── graceful-fs@4.1.11
  ├─┬ minimatch@3.0.4
  │ └─┬ brace-expansion@1.1.8
  │   ├── balanced-match@1.0.0
  │   └── concat-map@0.0.1
  ├─┬ readable-stream@2.3.3
  │ ├── core-util-is@1.0.2
  │ ├── isarray@1.0.0
  │ ├── process-nextick-args@1.0.7
  │ ├── safe-buffer@5.1.1
  │ ├── string_decoder@1.0.3
  │ └── util-deprecate@1.0.2
  └── set-immediate-shim@1.0.1
├── cli-table-redemption@1.0.1
├── commander@2.11.0
├─┬ cron@1.2.1
│ └── moment-timezone@0.5.13
├─┬ debug@3.1.0
│ └── ms@2.0.0
├── eventemitter2@1.0.5
├── fclone@1.0.11
├── UNMET OPTIONAL DEPENDENCY gkt@https://tgz.pm2.io/gkt-1.0.0.tgz
├─┬ mkdirp@0.5.1
│ └── minimist@0.0.8
├── moment@2.19.0
├─┬ needle@1.6.0
│ ├── debug@2.6.9
│ └── iconv-lite@0.4.19
├─┬ nssocket@0.6.0
│ ├── eventemitter2@0.4.14
│ └── lazy@1.0.11
├── pidusage@1.1.6
├─┬ pm2-axon@3.1.0
│ ├── amp@0.3.1
│ ├── amp-message@0.1.2
│ └── escape-regexp@0.0.1
```

```
├── pm2-axon-rpc@0.5.0
├── pm2-deploy@0.3.8
│   ├── async@1.5.2
│   └── tv4@1.3.0
├── pm2-multimeter@0.1.2
│   └── charm@0.1.2
├── pmx@1.5.4
│   ├── json-stringify-safe@5.0.1
│   └── vxx@1.2.2
│       ├── continuation-local-storage@3.2.0
│       │   ├── async-listener@0.6.8
│       │   │   └── emitter-listener@1.0.1
│       │   │       └── shimmer@1.0.0
│       │   ├── debug@2.6.9
│       │   ├── extend@3.0.1
│       │   ├── is@3.2.1
│       │   ├── lodash.findindex@4.6.0
│       │   ├── lodash.isequal@4.5.0
│       │   ├── lodash.merge@4.6.0
│       │   ├── methods@1.1.2
│       │   ├── shimmer@1.1.0
│       │   └── uuid@3.1.0
│       ├── promptly@2.2.0
│       │   ├── read@1.0.7
│       │   │   └── mute-stream@0.0.7
│       │   └── semver@5.4.1
│       ├── shelljs@0.7.8
│       │   ├── glob@7.1.2
│       │   │   ├── fs.realpath@1.0.0
│       │   │   ├── inflight@1.0.6
│       │   │   │   └── wrappy@1.0.2
│       │   │   └── once@1.4.0
│       │   ├── interpret@1.0.4
│       │   └── rechoir@0.6.2
│       │       └── resolve@1.4.0
│       │           └── path-parse@1.0.5
│       └── source-map-support@0.4.18
│           └── source-map@0.5.7
├── sprintf-js@1.1.1
├── vizion@0.2.13
│   └── async@1.5.2
└── yamljs@0.3.0
    ├── argparse@1.0.9
    └── sprintf-js@1.0.3
```

$ cat user.txt

/home/identv

```
└─┬ pm2-logrotate@2.3.0
  ├── graceful-fs@4.1.11
  ├─┬ moment-timezone@0.5.14
  │ └── moment@2.19.1
  ├─┬ node-schedule@1.2.5
  │ ├─┬ cron-parser@2.4.3
  │ │ └─┬ is-nan@1.2.1
  │ │   └─┬ define-properties@1.1.2
  │ │     ├── foreach@2.0.5
  │ │     └── object-keys@1.0.11
  │ ├── long-timeout@0.1.1
  │ └── sorted-array-functions@1.0.0
  └─┬ pm2@2.7.2
    ├─┬ async@2.5.0
    │ └── lodash@4.17.4
    ├── blessed@0.1.81
    ├─┬ chalk@1.1.3
    │ ├── ansi-styles@2.2.1
    │ ├── escape-string-regexp@1.0.5
    │ ├─┬ has-ansi@2.0.0
    │ │ └── ansi-regex@2.1.1
    │ ├── strip-ansi@3.0.1
    │ └── supports-color@2.0.0
    ├─┬ chokidar@1.7.0
    │ ├─┬ anymatch@1.3.2
    │ │ └─┬ micromatch@2.3.11
    │ │   ├─┬ arr-diff@2.0.0
    │ │   │ └── arr-flatten@1.1.0
    │ │   ├── array-unique@0.2.1
    │ │   ├─┬ braces@1.8.5
    │ │   │ ├─┬ expand-range@1.8.2
    │ │   │ │ └─┬ fill-range@2.2.3
    │ │   │ │   ├── is-number@2.1.0
    │ │   │ │   ├── isobject@2.1.0
    │ │   │ │   ├─┬ randomatic@1.1.7
    │ │   │ │   │ ├─┬ is-number@3.0.0
    │ │   │ │   │ │ └── kind-of@3.2.2
    │ │   │ │   │ └── kind-of@4.0.0
    │ │   │ │   └── repeat-string@1.6.1
    │ │   │ ├── preserve@0.2.0
    │ │   │ └── repeat-element@1.1.2
    │ │   ├─┬ expand-brackets@0.1.5
    │ │   │ └── is-posix-bracket@0.1.1
    │ │   └── extglob@0.3.2
```

```
│   │   │   │   ├── filename-regex@2.0.1
│   │   │   │   ├── kind-of@3.2.2
│   │   │   │   │   └── is-buffer@1.1.6
│   │   │   │   ├── object.omit@2.0.1
│   │   │   │   │   ├── for-own@0.1.5
│   │   │   │   │   │   └── for-in@1.0.2
│   │   │   │   │   └── is-extendable@0.1.1
│   │   │   │   ├── parse-glob@3.0.4
│   │   │   │   │   ├── glob-base@0.3.0
│   │   │   │   │   └── is-dotfile@1.0.3
│   │   │   │   ├── regex-cache@0.4.4
│   │   │   │   │   ├── is-equal-shallow@0.1.3
│   │   │   │   │   │   └── is-primitive@2.0.0
│   │   │   ├── normalize-path@2.1.1
│   │   │   │   └── remove-trailing-separator@1.1.0
│   │   ├── async-each@1.0.1
│   │   ├── UNMET OPTIONAL DEPENDENCY fsevents@^1.0.0
│   │   ├── glob-parent@2.0.0
│   │   ├── inherits@2.0.3
│   │   ├── is-binary-path@1.0.1
│   │   │   └── binary-extensions@1.10.0
│   │   ├── is-glob@2.0.1
│   │   │   └── is-extglob@1.0.0
│   │   ├── path-is-absolute@1.0.1
│   │   ├── readdirp@2.1.0
│   │   │   ├── minimatch@3.0.4
│   │   │   │   └── brace-expansion@1.1.8
│   │   │   │   │   ├── balanced-match@1.0.0
│   │   │   │   │   └── concat-map@0.0.1
│   │   │   ├── readable-stream@2.3.3
│   │   │   │   ├── core-util-is@1.0.2
│   │   │   │   ├── isarray@1.0.0
│   │   │   │   ├── process-nextick-args@1.0.7
│   │   │   │   ├── safe-buffer@5.1.1
│   │   │   │   ├── string_decoder@1.0.3
│   │   │   │   └── util-deprecate@1.0.2
│   │   │   └── set-immediate-shim@1.0.1
│   ├── cli-table-redemption@1.0.1
│   ├── commander@2.11.0
│   ├── cron@1.3.0
│   ├── debug@3.1.0
│   │   └── ms@2.0.0
│   ├── eventemitter2@1.0.5
│   ├── fclone@1.0.11
│   └── UNMET OPTIONAL DEPENDENCY gkt@https://tgz.pm2.io/gkt-1.0.0.tgz
```

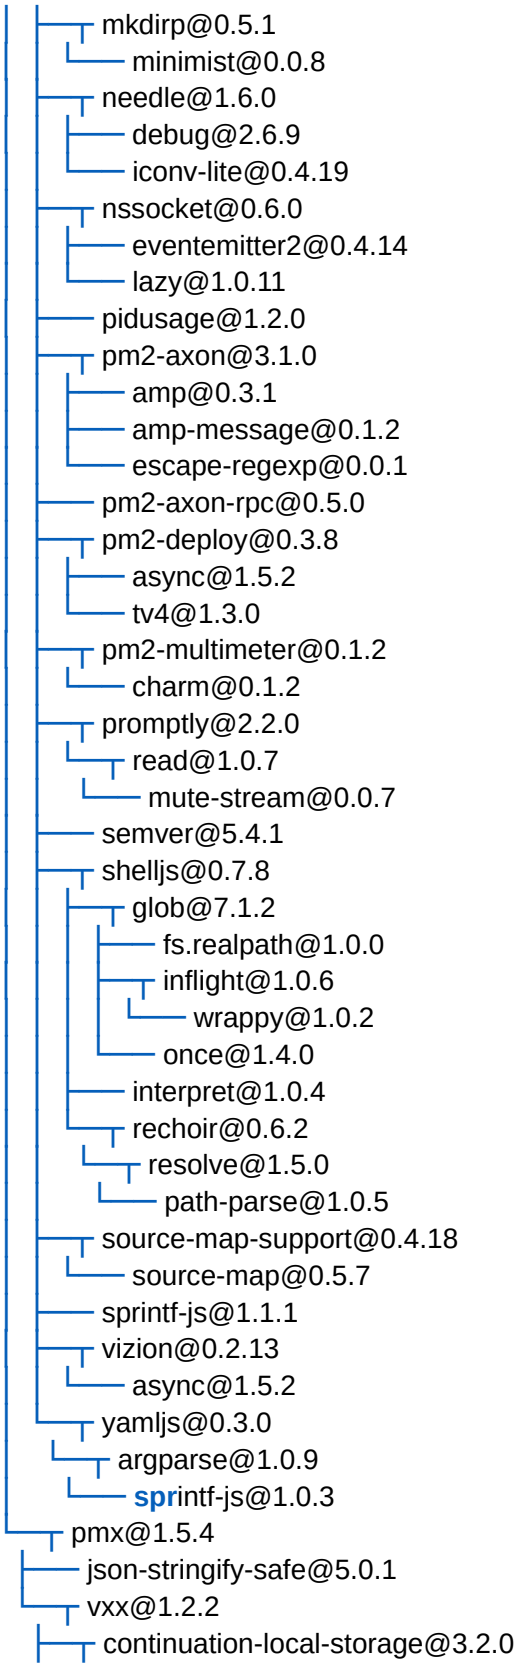
mkdirp@0.5.1
minimist@0.0.8
needle@1.6.0
debug@2.6.9
iconv-lite@0.4.19
nssocket@0.6.0
eventemitter2@0.4.14
lazy@1.0.11
pidusage@1.2.0
pm2-axon@3.1.0
amp@0.3.1
amp-message@0.1.2
escape-regexp@0.0.1
pm2-axon-rpc@0.5.0
pm2-deploy@0.3.8
async@1.5.2
tv4@1.3.0
pm2-multimeter@0.1.2
charm@0.1.2
promptly@2.2.0
read@1.0.7
mute-stream@0.0.7
semver@5.4.1
shelljs@0.7.8
glob@7.1.2
fs.realpath@1.0.0
inflight@1.0.6
wrappy@1.0.2
once@1.4.0
interpret@1.0.4
rechoir@0.6.2
resolve@1.5.0
path-parse@1.0.5
source-map-support@0.4.18
source-map@0.5.7
sprintf-js@1.1.1
vizion@0.2.13
async@1.5.2
yamljs@0.3.0
argparse@1.0.9
sprintf-js@1.0.3
pmx@1.5.4
json-stringify-safe@5.0.1
vxx@1.2.2
continuation-local-storage@3.2.0

- async-listener@0.6.8
  - emitter-listener@1.0.1
    - shimmer@1.0.0
- debug@2.6.9
- extend@3.0.1
- is@3.2.1
- lodash.findindex@4.6.0
- lodash.isequal@4.5.0
- lodash.merge@4.6.0
- methods@1.1.2
- shimmer@1.1.0
- uuid@3.1.0

Bower Mirror

Note

I'm even more unsure about this; I think it's basically another mirror for node.js packages. I believe this is a complete list of bower-managed packages installed for the tagger (like npm the tagger is the only part of the application that users bower):

$ bower list

bower check-new Checking for new versions of the project dependencies...
client#0.0.0 /home/foo/git/object-tagging/server/client

- angular-rangeslider#0.0.14
  - angular#1.5.11 (1.6.7-build.5496+sha.181ac0b available)
  - jquery#2.2.4 (3.2.1 available)
- angular-resource#1.5.11 (latest is 1.6.7-build.5496+sha.181ac0b)
  - angular#1.5.11 (latest is 1.6.7-build.5496+sha.181ac0b)
- angular-route#1.5.11 (latest is 1.6.7-build.5496+sha.181ac0b)
  - angular#1.5.11
- angular-sanitize#1.5.11 (latest is 1.6.7-build.5496+sha.181ac0b)
- angular-ui-bootstrap#1.3.3 (latest is 2.5.6)
- angular-ui-router#0.3.2 (latest is 1.0.10)
  - angular#1.5.11 (1.6.7-build.5496+sha.181ac0b available)
- angularjs#1.5.11 (latest is 1.6.7-build.5496+sha.181ac0b)
- bootstrap#3.3.7 (latest is 4.0.0-beta.2)
  - jquery#2.2.4 (3.2.1 available)
- bootstrap-sass-official#3.3.7
  - jquery#2.2.4
- fabric.js#1.7.19 (latest is 2.0.0-beta.7)
- font-awesome#4.6.3 (latest is 4.7.0)
- jquery#2.2.4 (latest is 3.2.1)
- jsfeat#0.0.8

- moment#2.19.0 (2.19.1 available)
- ng-file-upload#12.2.13
  - angular#1.5.11 (1.6.7-build.5496+sha.181ac0b available)
- stormpath-sdk-angularjs#1.1.1 (latest is 2.0.1)
  - angular#1.5.11 (1.6.7-build.5496+sha.181ac0b available)
- tracking#1.1.3
  - dat-gui#0.5.0 (latest is 0.6.5)
  - threejs#r67
- underscore#1.8.3
- videogular#1.4.4
  - angular#1.5.11 (1.6.7-build.5496+sha.181ac0b available)
  - angular-sanitize#1.5.11 (1.6.7-build.5496+sha.181ac0b available)
    - angular#1.5.11
- videogular-buffering#1.4.4
  - videogular#1.4.4
- videogular-controls#1.4.4
  - videogular#1.4.4
- videogular-dash#1.4.4
  - videogular#1.4.4
- videogular-overlay-play#1.4.4
  - videogular#1.4.4
- videogular-poster#1.4.4
  - videogular#1.4.4
- videogular-themes-default#1.4.4
  - videogular#1.4.4

12. Table of Figures

13. Table of Document Tables