



Correlating Autopsy Cases

Brian Carrier

OSDFCon 2017

Imagine This World.....

- You get a new piece of media and are quickly alerted because it has a phone number that was also in a case from a month ago.
- You find a document with intelligence and you want to know if it has been seen before. You quickly determine that it was also on media from John Doe.
- The suspect says the child pornography was placed there by malware. You can quickly show that it was on multiple devices.

It's Now Possible



If you aren't a 5-year old....



Before The 4.5.0 Release

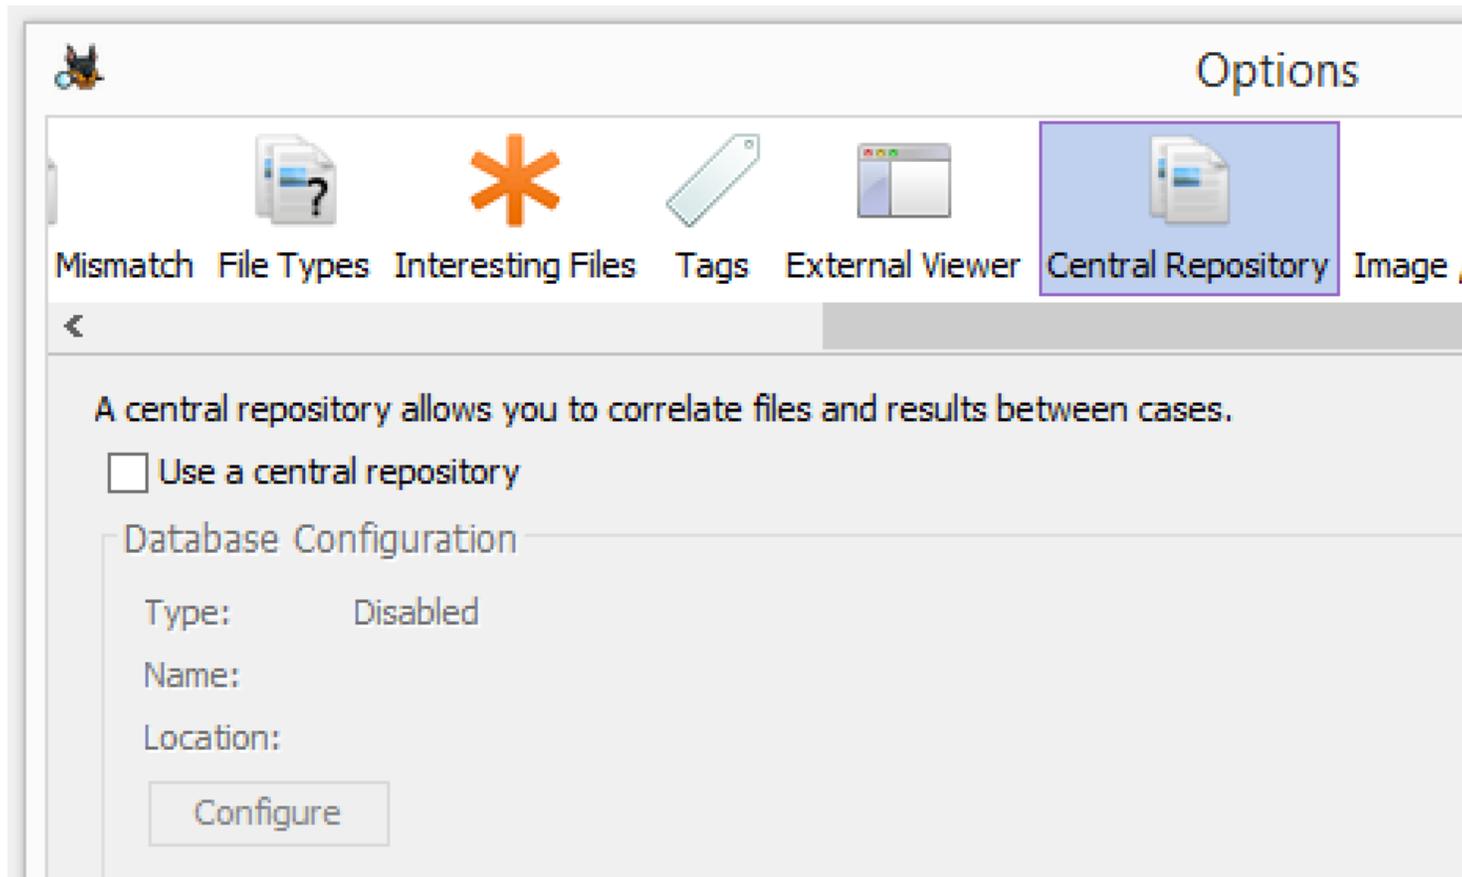
- When you made a case, Autopsy created a database.
 - Single-user Cases: A SQLite database in the case folder.
 - Multi-user Cases: A new database on the PostgreSQL server.
- The database contains:
 - File system information (file metadata and names, partitions, etc.)
 - The Blackboard (web bookmarks, keyword hits, etc.)
- Does not contain file content or any data that spanned cases.
- This makes it easy to scale because the databases stay small.

Now (as of 4.5.0)

- Autopsy still maintains a single database per case.
- It can now maintain a non-case-specific database.
 - We call it the central repository.
- Can be used for:
 - Correlation
 - Hash Databases
 -

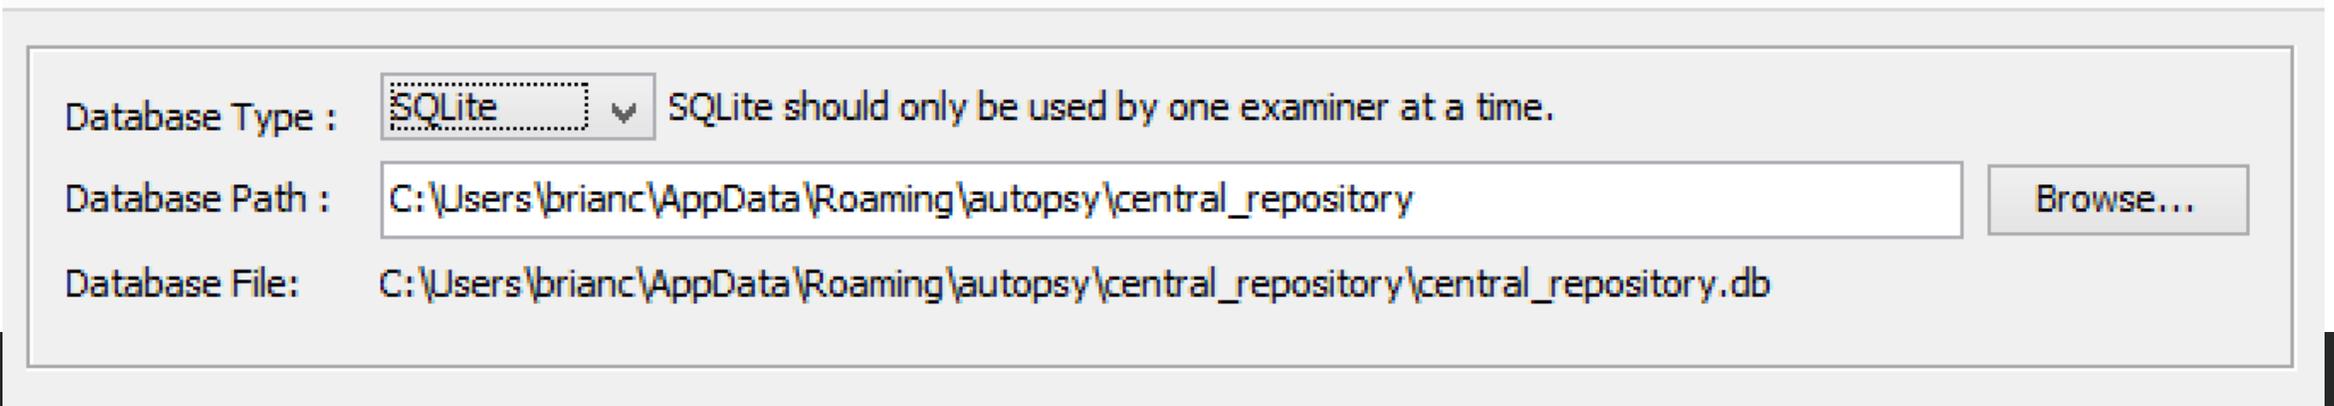
Configuration: Enabling It

- Step 1: Enable It!
 - Tools -> Options -> Central Repository panel.



Configuration: Database Type

- Two types are supported.
- SQLite:
 - Database is stored in a folder.
 - Default location is in AppData.
 - Requires no other installations.
 - BUT, can be used by only one user at a time. Do not put on a network share and have multiple examiners using it at the same time.



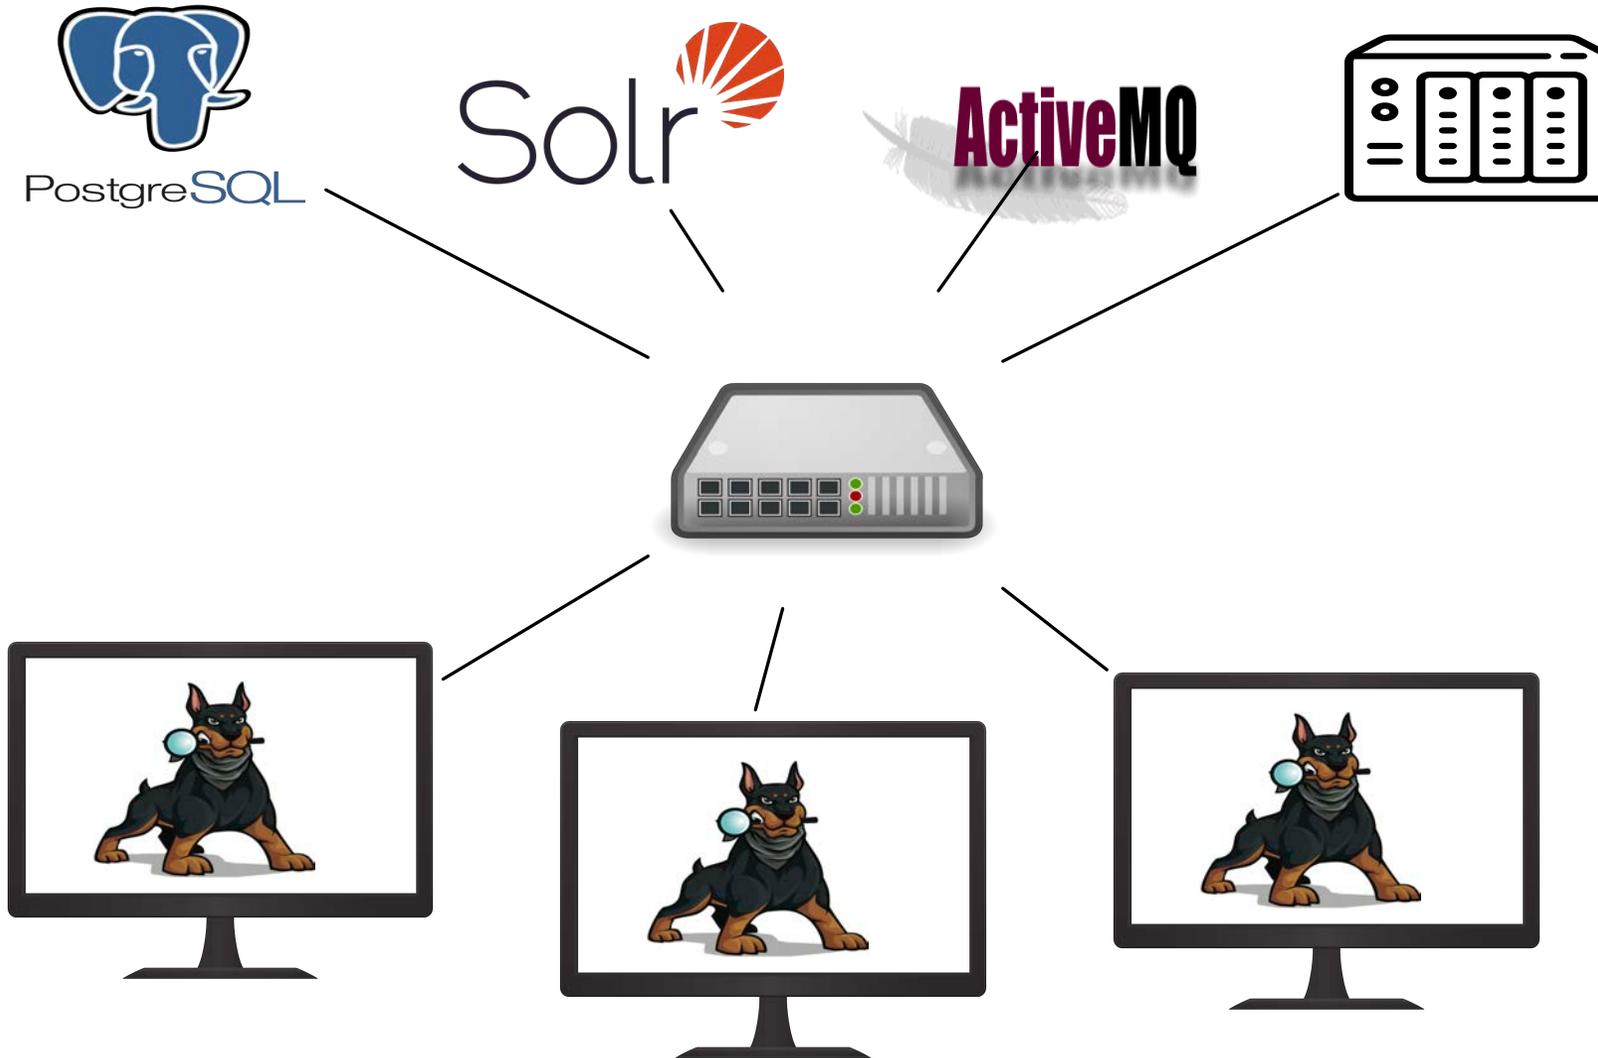
The screenshot shows a configuration dialog box with the following fields and text:

- Database Type :** A dropdown menu with "SQLite" selected. To its right, a warning message reads: "SQLite should only be used by one examiner at a time."
- Database Path :** A text input field containing the path "C:\Users\brianc\AppData\Roaming\autopsy\central_repository". To its right is a "Browse..." button.
- Database File:** A text input field containing the file path "C:\Users\brianc\AppData\Roaming\autopsy\central_repository\central_repository.db".

- PostgreSQL
 - Database is stored on a server
 - Can be used by multiple users at a time.
 - You must install and configure the PostgreSQL server.
 - Can use the same server for multi-user cases.

Database Type :	<input type="text" value="PostgreSQL"/>
Host Name / IP :	<input type="text" value="db-forensics-server"/>
Port :	<input type="text" value="5432"/>
User Name :	<input type="text" value="Database User"/>
User Password :	<input type="password"/>

Multi-user Cases



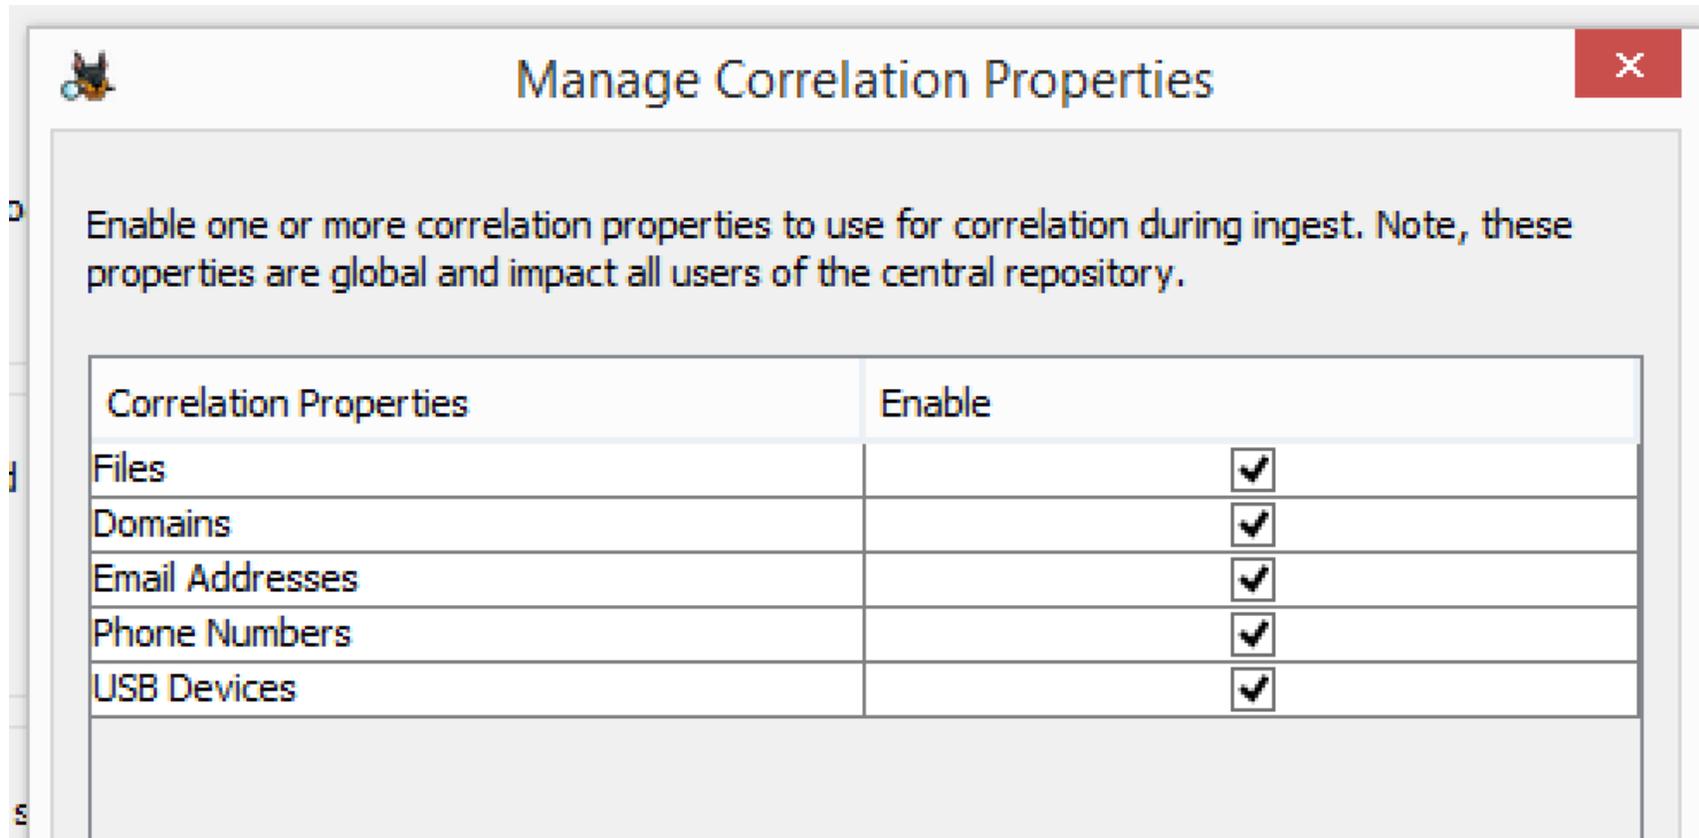
- If you are a single-person shop, stick with SQLite.
- If there are multiple people in your lab, setup PostgreSQL.
 - It's fairly easy.
 - Follow the instructions in the Autopsy docs.

Now What? Basic Correlation

- What: Allows you to find links with previous cases
- How:
 - Information about each file, phone number, etc. is stored in the central repo when a data source is “ingested”.
 - When you select an item, you can see its other occurrences.

Correlation Setup

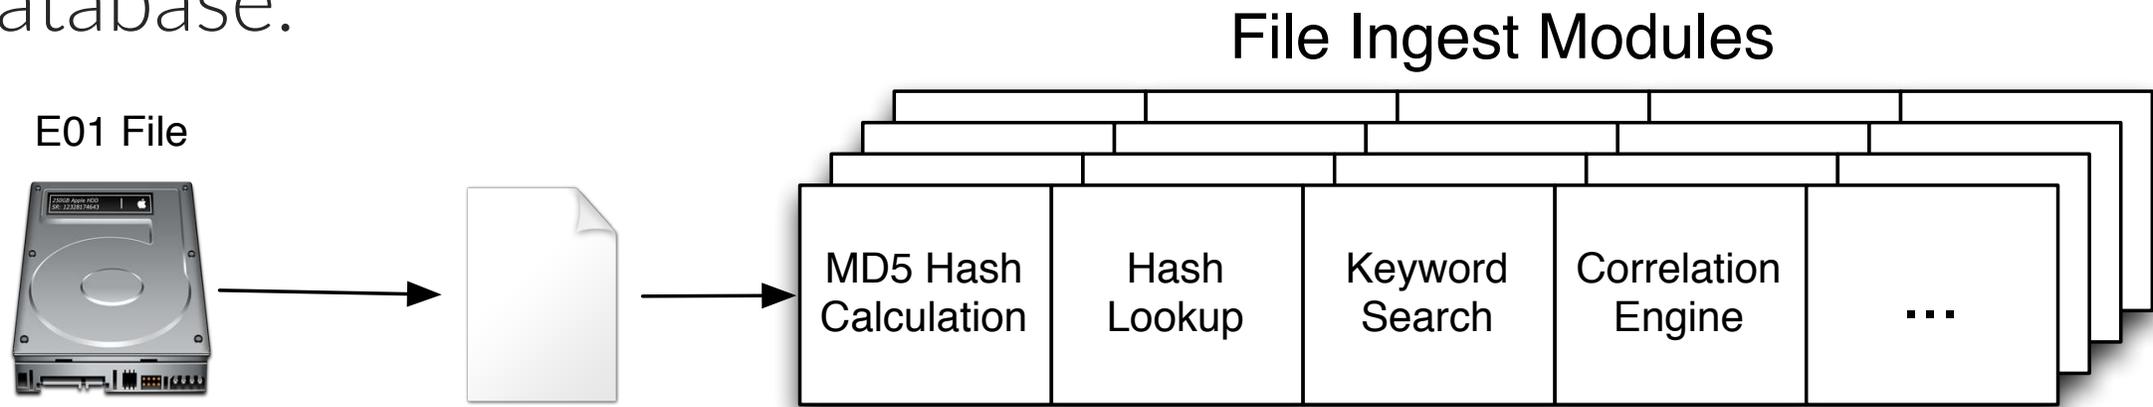
- You can refine what types of “properties” to correlate.
- In the Options panel.



- File: The MD5 and path for each file.
- Domain: From web artifact URLs and keyword hits.
- Email Addresses: From email messages and keyword hits.
- Phone Numbers: From messages, contact books, call logs, and keyword hits.
- USB Devices: From the devices plugged in (based on the registry).

Properties are Saved During Ingest

- Hash, Keyword Search, Email, etc. modules must be enabled to extract data.
- Correlation Engine ingest module will save that data to the database.



- NOTE: If you don't enable the initial modules, the data won't be saved.

Seeing Correlations: Basic Autopsy Layout

Close Case + Add Data Source Generate Report

Warning Keyword Lists Keyword Search

Directory Listing /LogicalFileSet1/zombies 5 Results

Table Thumbnail

Name	Location	Modified Time	Change Time
 1365273819-zombies-invade-brussels-streets	/LogicalFileSet1/zombies/1365273819-zombies-invade-brussels-street...	0000-00-00 00:00:00	0000-00-00 00:00:00
 d2z9afumccmca6ffe2zi.jpg	/LogicalFileSet1/zombies/d2z9afumccmca6ffe2zi.jpg	0000-00-00 00:00:00	0000-00-00 00:00:00
 walking-dead-zombie.jpg	/LogicalFileSet1/zombies/walking-dead-zombie.jpg	0000-00-00 00:00:00	0000-00-00 00:00:00
 walkingdead_ap.jpg	/LogicalFileSet1/zombies/walkingdead_ap.jpg	0000-00-00 00:00:00	0000-00-00 00:00:00
 zombie-apocalypse.jpg	/LogicalFileSet1/zombies/zombie-apocalypse.jpg	0000-00-00 00:00:00	0000-00-00 00:00:00

Hex Strings File Metadata Results Indexed Text Media



“Other Occurrences” Tab

- The lower right tab will show you which data sources this item also occurred in.
- If the selected item has:
 - A file associated with it, MD5 will be used.
 - An email (such as contact book or message), it will be searched.
 -
- Occurrences are shown both within the current case and other cases.

“Other Occurrences” Tabs

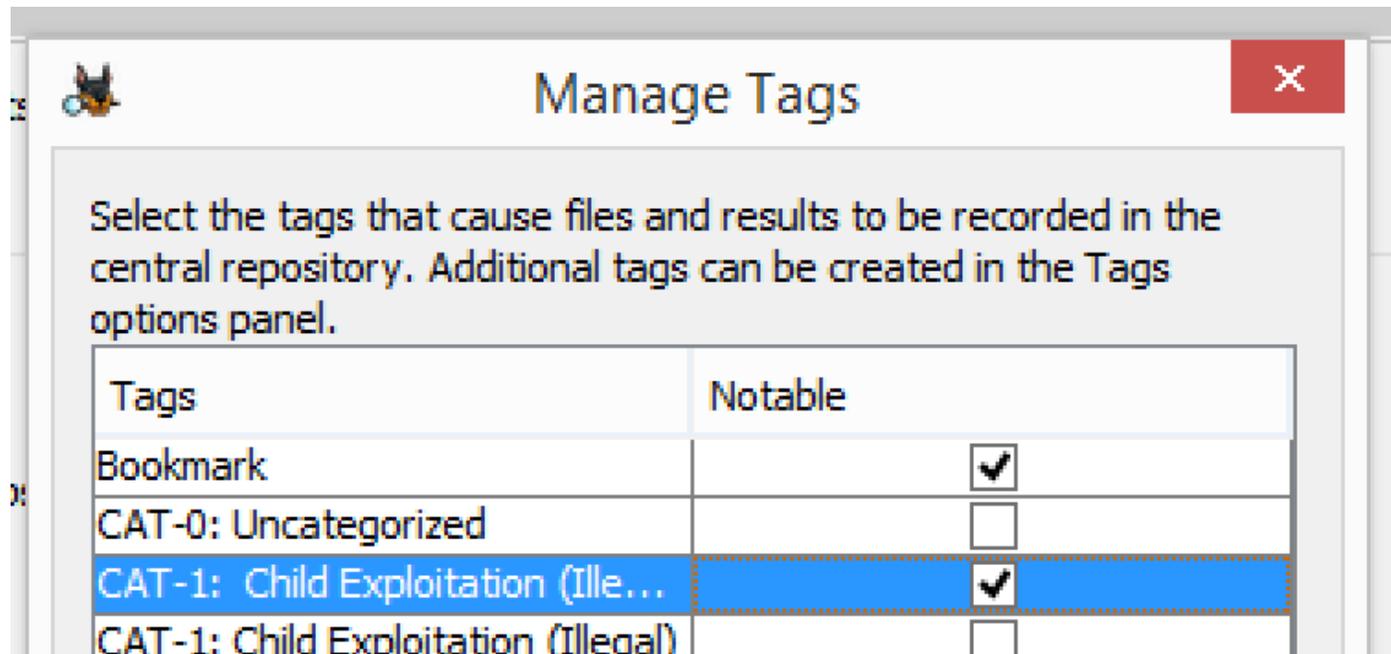
Hex	Strings	File Metadata	Results	Indexed Text	Media	Other Occurrences	Video Triage	Text Gist
Case	Data Source	Correlation Type	Correlation Value	Known	Scope	Path		
demo-case123d	xp-sp3-v3.001	Files	af1748c6894effd15e8a97a291d20357	unknown	Local	/documents and settings/john/local settings/tempora		

Now What? Previously “Bad” / Notable

- What: Allows you to see if a previous case considered an item to be notable.
- How:
 - When a user tags an item as notable, that gets saved in the central repository.
 - When it is seen again, it gets flagged.

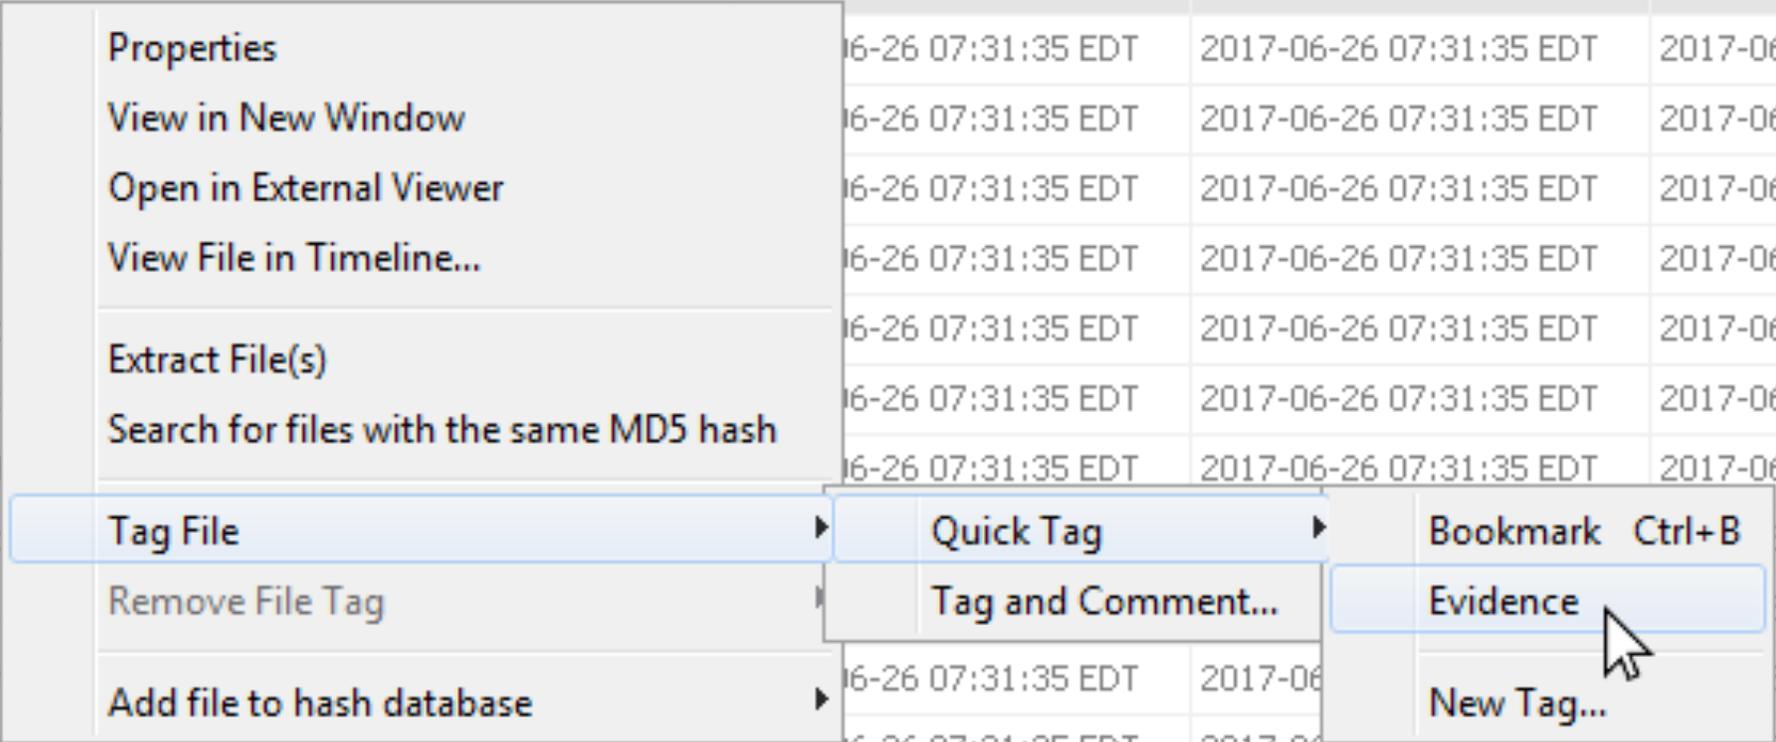
Previously Notable: Setup

- You need to configure which tag names are associated with “notable”.
- “Manage Tags” in the Options panel.



Previously Notable: Tagging a File

0000_d.txt	2017-06-22 20:16:30 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	11
0000_e.txt	2017-06-22 20:16:30 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	11
0000_f.txt	2017-06-22 20:16:30 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	11
0000_g.txt	2017-06-22 20:16:30 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	11
0000_h.txt	2017-06-22 20:16:30 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	11
0000_i.txt	2017-06-22 20:16:30 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	11
0000_j.txt	2017-06-22 20:16:30 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	11
0000_k.txt	2017-06-22 20:16:30 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	11
0000_l.txt	2017-06-22 20:16:30 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	11
0000_m.txt	2017-06-22 20:16:30 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	11
0000_n.txt	2017-06-22 20:16:30 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	11
0000_o.txt	2017-06-22 20:16:32 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	11
0000_p.txt	2017-06-22 20:16:32 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	11
0000_q.txt	2017-06-22 20:16:32 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	2017-06-26 07:31:35 EDT	11



- The instance of the item gets stored as being “notable”
 - An instance is an occurrence at a specific ‘path’ and ‘case’.
- NOTE: There could be other occurrences of that file, email, etc. They will not be marked as “notable”.
- If you untag the file, its “notable” status will be removed from the central repository.

- Enable the Correlation Engine ingest module (just like for the correlation feature).
- It will query the central repository for previous notable occurrences of the item.
- If any are found, an “Interesting Item” artifact will be created.
- You can find it in the tree and an inbox message will be created.

Previously Notable: Seeing Results

- Data Sources
- Views
- Results
 - Extracted Content
 - Keyword Hits
 - Hashset Hits
 - E-Mail Messages
 - Interesting Items
 - Previously Tagged As Notable (Central Repository) (1)
 - Interesting Files (1)
 - Interesting Results (0)
 - Accounts
- Tags
- Reports

Table Thumbnail		
Source File	Comment	File Path
 image_normal[1].jpg	Previous Case: demo-33323	/img_xp-sp3

<

Hex Strings File Metadata Results Indexed Text Media Other Occurrences



The Future: Tighter Integration

- The January release (4.6.0) will use the repository for hash databases.
- NSRL and notable hashsets can be shared in multi-user cases.
- Users will be able to:
 - Import and create hashsets into the central repo
 - Pick which hashsets to use

Hash Configuration Same as Local DBs

Run ingest modules on:

All Files, Directories, and Unallocated Space

- Recent Activity
- Hash Lookup
- File Type Identification
- Embedded File Extractor
- Exif Parser
- Keyword Search
- Email Parser
- Extension Mismatch Detector
- E01 Verifier
- Interesting Files Identifier

Select known hash databases to use:

- NSRFile.txt-255m-md5

Select notable hash databases to use:

- notable_hash_db.txt-md5

- With 4.5.0, you need to configure tags as 'notable' separate from where you define tags.
- The 4.6.0 release will force you to decide what a tag name means when you create it.

- Central Repository allows for more complex analytics.
- See if an item has been seen before
- See if an item has been previously marked as notable
- Easier hash database management

- It's all free and open source....

<http://sleuthkit.org/autopsy/>

Questions

Brian Carrier

brianc@basistech.com

- Tables for:
 - Case information
 - Data source information
 - Each type of correlation property
- Correlation Properties Tables:
 - Case and data source identifiers
 - Value (MD5, email address, etc.)
 - Path
 - Known status (notable, etc.)
 - Comment