

# Heap Spray Detection with Heap Inspector

---

Blackhat USA 2011

**Aaron LeMasters**

**8/3/2011**

Heap Inspector is a heap visualization and analysis tool. It has the ability to collect a process's heaps using both API and raw methods. Features include searching heaps for string or byte patterns (including regex), dumping heap chunks to a file, and viewing chunks in a hex editor pane. Heaps are displayed visually in a bar chart format known as the heap hash map, allowing the user to view allocations spatially. A similar chart called the heap data map overlays regular expression matches for useful patterns on top of the heap bars. This visualization allows an investigator to quickly discover evidence of a heap spray attack and other useful information stored in an application's heap memory.



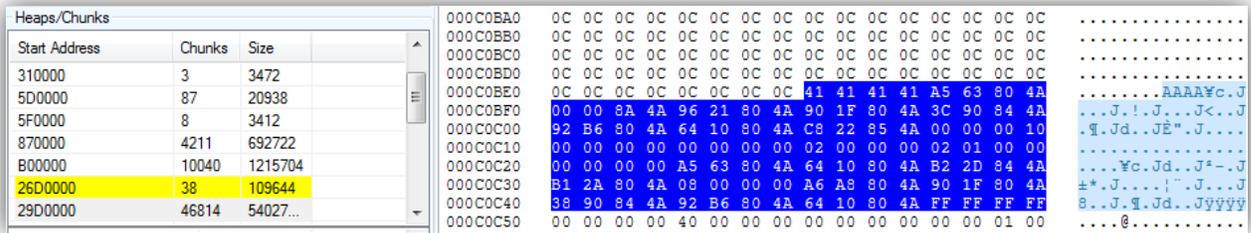


Figure 2: Extracting the shellcode in heap memory from an exploit of CVE-2010-2883

Once extracted, the shellcode can be analyzed statically in a disassembler or dynamically in an emulator to determine its exact functionality. In summary, the heap hash map makes it trivial to determine the trustworthiness of any suspect document (PDF, MS Word, Excel, etc) simply by running it in a sandbox or virtual machine.

### Searching for Personally Identifiable Information

Similar to the heap hash map, the *heap data map* is used to visualize heap allocations, but in this case only heap chunks that match a supplied list of regular expression are drawn on the map. Several default regular expressions are included with Heap Inspector, such as File/UNC Path, GUID, IPv4 address, SQL queries, URL/domains, Social security numbers, and credit card numbers. You can create, import and export your own regular expressions. As shown below, only regular expressions that are matched in heap memory will be displayed.




Figure 3: A portion of the heap data map for Microsoft Outlook

In addition to the graphical heap data map, the search tab provides a simple interface to search for regular expressions in application heap memory. The screenshot below shows an active conversation found in Skype’s heap memory. There are many other useful items stored in Skype’s memory, such as Delphi executable files, SQL statements, and various URLs and configuration data. As a side note, since Heap Inspector reads memory that’s in use by the application, anti-debugging and obfuscation tricks do not work (Skype deploys various anti-debugging tricks to prevent attaching with a debugger).



server DLL uses either Win32 API or raw methods to build a list of available *heaps* and *chunks* (contained in *segments*). The heap chunks contain the actual data we want to collect and analyze. For an in-depth discussion of Windows heap internals, see Chris Valacek's talk [1]. Because of the way Heap Inspector works, it must be executed with Administrator credentials.

## Caveats and Technical challenges

Since Heap Inspector relies on DLL injection to collect heap information, it is subject to all of the caveats and technical restrictions inherent to that technique. This includes the possibility of causing instability in the target application as a result of synchronization issues [2]. Such issues arise by virtue of adding a new thread to a remote process that might not be thread-safe. Additionally, accessing and/or locking private heaps might cause deadlock conditions. Furthermore, when the DLL is injected (not using the reflective method), the entry point to each DLL in the process is called. This might result in abnormal or undefined behavior.

There are some tricks application developers can employ to prevent standard DLL injection (e.g., `VirtualAlloc`, `WriteProcessMemory`, `CreateRemoteThread`), such as IAT function hooking and least-privilege concepts (e.g., the process runs with minimal privileges, in a restricted job object, and on its own desktop). Google Chrome uses all of these techniques, making injection difficult. In the case of Google Chrome and other sandboxed processes, even if the DLL is successfully injected using Reflective DLL Injection [3], the injected code can't do anything without first escalating privileges. Due to these obstacles, Heap Inspector is not compatible with Google Chrome or sandboxed processes. Note that it is possible to inject into Chrome's broker process and plugin processes, but this is generally not very useful.

There are other situations where injection might fail, such as across session boundaries in Terminal Services or Vista session separation and in System/service processes. These issues can usually be overcome by injecting as a service or impersonating a user to switch sessions.

When injecting, you must consider the target platforms (32-bit, 64-bit and 32-bit emulated WoW64). There are a few ways to successfully inject across all platforms, but the method used by Heap Inspector is to have an injector stub compiled for each supported platform (32-bit and 64-bit). The injector stub relies on a parameter structure written to the remote process that contains a table of `kernel32` function addresses. For 32-bit and 64-bit native platforms, the .NET application can collect these addresses itself. To support WoW64 processes, the addresses are collected by launching a separate 32-bit application (`Stub32.exe`).

Other miscellaneous challenges include:

- For exception handling in the injected DLL, use OS-provided SEH instead of C++ exception handling, as you have first-chance handling and more fine-grained control over the exception.
- The `smss.exe` process doesn't fully map in `kernel32.dll`, so you cannot inject into this process with Heap Inspector. If you attempted to do so, an access violation would occur, which causes a BSOD.

## How is Heap Inspector different than other tools?

There are many ways to collect data stored in memory, such as attaching with a debugger or dumping physical memory using a host of available tools. The most powerful and interactive of these options is a debugger such as `Windbg`, which requires technical expertise and cryptic low-level commands. Physical memory acquisition and analysis is powerful but also requires technical expertise and knowledge of memory structures and operating system internals. It also has the disadvantage of being a “smear” of the contents in memory and instantly stale once acquired.

In general, Heap Inspector seeks a middle ground between technical savvy and functionality. It is as close to “real time” as you can get without using a debugger and does not require the technical savvy to navigate memory contents. Typically in memory analysis, an investigator is dealing with either an entire process dump or an entire physical memory dump, which can be an overwhelming amount of data to cull through. By collecting data only from an application’s heap, an investigator can focus on data that’s being actively allocated and used by the program.

With regards to string searching functionality, most investigators are familiar with tools such as SysInternals’ `strings` utility, which will dump all strings in a binary. Similarly, many memory analysis tools allow an investigator to collect all strings in a process address space, which will include strings from the process’s image. In contrast, Heap Inspector only collects strings from the heap, which represents dynamically allocated data used through the life of the program. This distinction is important because it eliminates a large portion of data that might not be useful.

Finally, when applied to heap spray detection, Heap Inspector uses a simple CRC32 frequency-of-occurrence concept to detect and visualize a successful heap spray. All of the current research and tools available for heap spray detection and prevention [ 5, 6, 7 ] use a combination of binary instrumentation, emulation, hooking and/or dynamic code analysis to make an informed decision about the contents of heap data. While this has the added benefit of detecting and preventing heap spray attacks in “real time”, it has the notable drawbacks of adding overhead and generating false positives/negatives.

## Conclusion

Heap Inspector is an easy-to-use, free tool that helps an investigator visualize heap data. Uses of this tool include post-mortem heap spray detection, auditing of personally-identifiable information in application heap data, and reversing of proprietary memory structures stored on the heap.

## Future Direction

Currently, Heap Inspector only supports live system analysis by injecting into running processes. This method uses API techniques. In future versions of the tool, raw methods to parse heap structures by using the `PEB` will be employed so that offline memory images can be analyzed.

Also, since the heap spray detection is post-mortem, Heap Inspector will only detect the heap spray after the attack was successful. A future release of the tool might include the capability to detect such

an attack in progress. However, several tools and research already exist for this purpose [4, 5, 6, 7, 8] and rely on some variations of binary instrumentation and dynamic code analysis.

### Further reading/references

[1] Chris Valasek, *Understanding the Low Fragmentation Heap*, [http://illmatics.com/Understanding\\_the\\_LFH.pdf](http://illmatics.com/Understanding_the_LFH.pdf)

[2] Microsoft, *CreateRemoteThread*, <http://msdn.microsoft.com/en-us/library/ms682437%28VS.85%29.aspx>

[3] Stephen Fewer, *Reflective DLL Injection*, [http://www.harmonysecurity.com/files/HS-P005\\_ReflectiveDllInjection.pdf](http://www.harmonysecurity.com/files/HS-P005_ReflectiveDllInjection.pdf)

[4] Didier Stevens, *HeapLocker*, <http://blog.didierstevens.com/2010/12/06/heaplocker/>

[5] Microsoft Research, *Nozzle*, <http://research.microsoft.com/apps/pubs/default.aspx?id=76528>

[6] Microsoft Research, *Enhanced Mitigation Experience Toolkit v2.0*, <http://www.microsoft.com/download/en/details.aspx?id=5419>

[7] Akritidis, et al, *STRIDE: Polymorphic Sled Detection Through Instruction Sequence Analysis*, <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.79.5094&rep=rep1&type=pdf>

[8] Bania, *JIT Spraying and Mitigations*, [http://www.kryptoslogic.com/download/JIT\\_Mitigations.pdf](http://www.kryptoslogic.com/download/JIT_Mitigations.pdf)