

[\(https://breakingmalware.com/\)](https://breakingmalware.com/)

[\(https://breakingmalware.com/\)](https://breakingmalware.com/)



[\(https://breakingmalware.com/\)](https://breakingmalware.com/)

[\(https://breakingmalware.com/\)](https://breakingmalware.com/) Sponsored by enSilo (<http://www.ensilo.com/>)



- [Tal Liberman \(https://breakingmalware.com/author/tal-liberman/\)](https://breakingmalware.com/author/tal-liberman/)
- [October 27, 2016 \(https://breakingmalware.com/injection-techniques/atombombing-brand-new-code-injection-for-windows/\)](https://breakingmalware.com/injection-techniques/atombombing-brand-new-code-injection-for-windows/)
- [49 Comments \(https://breakingmalware.com/injection-techniques/atombombing-brand-new-code-injection-for-windows/#comments\)](https://breakingmalware.com/injection-techniques/atombombing-brand-new-code-injection-for-windows/#comments)

## **[AtomBombing: Brand New Code Injection for Windows](https://breakingmalware.com/injection-techniques/atombombing-brand-new-code-injection-for-windows/)** **[\(https://breakingmalware.com/injection-techniques/atombombing-brand-new-code-injection-for-windows/\)](https://breakingmalware.com/injection-techniques/atombombing-brand-new-code-injection-for-windows/)**

**TL;DR** Here's a new code injection technique, dubbed AtomBombing, which exploits Windows atom tables and Async Procedure Calls (APC). Currently, this technique goes undetected by common security solutions that focus on preventing infiltration.

Code injection has been a strong weapon in the hacker's arsenal for many years. For background on code injection and its various uses in APT type attack scenarios please take a look at:

<http://blog.ensilo.com/atombombing-a-code-injection-that-bypasses-current-security-solutions> (<http://blog.ensilo.com/atombombing-a-code-injection-that-bypasses-current-security-solutions>)

### **Overview**

I started poking around to see how hard it would be for a threat actor to find a new method that security vendors are unaware of and bypasses most security products. It also needed to work on different processes rather than being tailored to fit a specific process.

I would like to introduce you to AtomBombing – a brand new code injection technique for Windows.

AtomBombing works in three main stages:

1. Write-What-Where – Writing arbitrary data to arbitrary locations in the target process's address space.
2. Execution – Hijacking a thread of the target process to execute the code written in stage 1.

3. Restoration – Cleaning up and restoring the execution of the thread hijacked in stage 2.

## AtomBombing Stage 1: Write-What-Where

I stumbled onto a couple of rather interesting API calls:

- **GlobalAddAtom** ([https://msdn.microsoft.com/en-us/library/windows/desktop/ms649060\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/desktop/ms649060(v=vs.85).aspx)) – Adds a character string to the global atom table and returns a unique value (an atom) identifying the string.
- **GlobalGetAtomName** ([https://msdn.microsoft.com/en-us/library/windows/desktop/ms649063\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/desktop/ms649063(v=vs.85).aspx)) – Retrieves a copy of the character string associated with the specified global atom.

By calling GlobalAddAtom one can store a null terminated buffer in the global atom table. This table is accessible from every other process on the system. The buffer can then be retrieved by calling GlobalGetAtomName. GlobalGetAtomName expects a pointer to an output buffer, so the caller chooses where the null terminated buffer will be stored.

In theory, if I could add a buffer containing shellcode to the global atom table by calling GlobalAddAtom, and then somehow get the target process to call GlobalGetAtomName I could copy code from my process to the target process, without calling WriteProcessMemory.

Calling GlobalAddAtom from my process is pretty straightforward, but how would I get the target process to call GlobalGetAtomName?

By using Async Procedure Calls (APC):

*QueueUserApc* – adds a user-mode asynchronous procedure call (APC) object to the APC queue of the specified thread.

```
DWORD WINAPI QueueUserAPC(  
_In_ PAPCFUNC pfnAPC,  
_In_ HANDLE hThread,  
_In_ ULONG_PTR dwData  
);
```

QueueUserApc receives a pointer to an APCProc which is defined as follows:

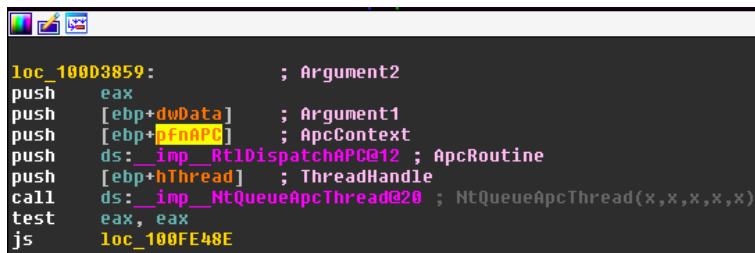
```
VOID CALLBACK APCProc(  
_In_ ULONG_PTR dwParam  
);
```

GlobalGetAtomName's prototype is:

```
UINT WINAPI GlobalGetAtomName(  
_In_ ATOM nAtom,  
_Out_ LPTSTR lpBuffer,  
_In_ int nSize  
);
```

Since GlobalGetAtomName expects 3 parameters (while APCProc is defined to expect only 1 parameter) we can't use QueueUserApc to get the target process to call GlobalGetAtomName.

Let's take a look at the internals of QueueUserApc:



```
loc_100D3859:          ; Argument2  
push    eax  
push    [ebp+dwData]  ; Argument1  
push    [ebp+pfnAPC]  ; ApcContext  
push    ds:imp_RtlDispatchAPC@12 ; ApcRoutine  
push    [ebp+hThread] ; ThreadHandle  
call    ds:imp_NtQueueApcThread@20 ; NtQueueApcThread(x,x,x,x,x)  
test    eax, eax  
js     loc_100FE48E
```

(./AtomBombing\_Brand New Code Injection for Windows - Breaking Malware\_files/Figure-0-QueueUserApc-Internals.png)

Figure 1: QueueUserApc

As you can see QueueUserApc uses the undocumented NtQueueApcThread syscall in order to add the APC to the target thread's APC queue.

Interestingly enough NtQueueApcThread receives a pointer to a function that is to be called asynchronously in the target thread, but the function being passed is not the original APCProc function the caller passed to QueueUserApc. Instead the function being passed is ntdll!RtlDispatchAPC, and the original APCProc function passed to QueueUserApc is passed as a parameter to ntdll!RtlDispatchAPC.

Let's take a look at ntdll!RtlDispatchAPC:

```

mov     esi, [ebp+p_ActivationContext]
cmp     esi, 0FFFFFFFh
jnz     loc_4830E788

; START OF FUNCTION CHUNK FOR _RtlDispatchAPC@12

loc_4830E788:
mov     edx, esi
lea     ecx, [ebp+var_40]
call   RtlActivateActivationContextUnsafeFast(x,x)
and     [ebp+ms_exc.registration.TryLevel], 0
push   [ebp+p_ApcProcParam]
mov     ecx, [ebp+p_ApcProc]
call   ds:__guard_check_icall_fptr
call   [ebp+p_ApcProc]
mov     [ebp+ms_exc.registration.TryLevel], 0FFFFFFFh
call   lblDeactivateActivationContext

lblDeactivateActivationContext:
lea     ecx, [ebp+var_40]
call   RtlDeactivateActivationContextUnsafeFast(x)
push   esi
call   RtlReleaseActivationContext(x)
retn

; END OF FUNCTION CHUNK FOR _RtlDispatchAPC@12

push   [ebp+p_ApcProcParam]
mov     ecx, [ebp+p_ApcProc]
call   ds:__guard_check_icall_fptr
call   [ebp+p_ApcProc]

loc_4830E7B1:
jmp     loc_4B2A464F

```

(./AtomBombing\_Brand New Code Injection for Windows - Breaking Malware files/Figure-1-RtlDispatchApc.png)

Figure 2: ntdll!RtlDispatchAPC

It starts by checking if the 3<sup>rd</sup> parameter is valid, which means an ActivationContext needs to be activated before dispatching the APC.

If an ActivationContext needs to be activated:

```

; START OF FUNCTION CHUNK FOR _RtlDispatchAPC@12

loc_4830E788:
mov     edx, esi
lea     ecx, [ebp+var_40]
call   RtlActivateActivationContextUnsafeFast(x,x)
and     [ebp+ms_exc.registration.TryLevel], 0
push   [ebp+p_ApcProcParam]
mov     ecx, [ebp+p_ApcProc]
call   ds:__guard_check_icall_fptr
call   [ebp+p_ApcProc]
mov     [ebp+ms_exc.registration.TryLevel], 0FFFFFFFh
call   lblDeactivateActivationContext

```

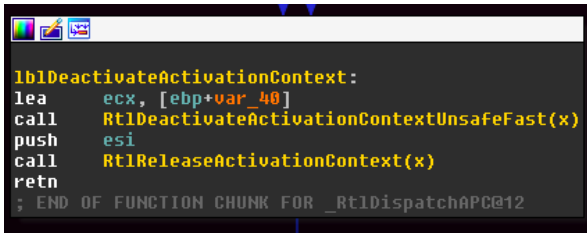
(./AtomBombing\_Brand New Code Injection for Windows - Breaking Malware files/Figure-2-RtlDispatchApc-2.png)

Figure 3: ntdll!RtlDispatchAPC – RtlActivateActivationContextUnsafeFast

The function ntdll!RtlDispatchAPC executes the following:

1. The passed ActivationContext (currently in ESI) will be activated by calling RtlActivateActivationContextUnsafeFast.
2. The parameter to the original APCProc function (i.e. the third parameter passed to QueueUserApc) is pushed onto the stack. This is because we are about to call the original APCProc function.
3. Right before dispatching the APC, a call to CFG (\_\_guard\_check\_icall\_fptr) is made to make sure the APC target is a CFG valid function.
4. A call to the original APCProc is made, and that's it – the APC has been dispatched.

Once APCProc returns, the activation context is deactivated:

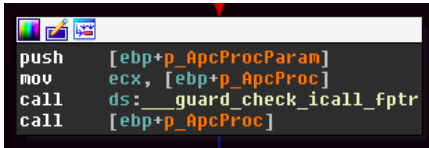


```
!b1DeactivateActivationContext:
lea     ecx, [ebp+var_40]
call   RtlDeactivateActivationContextUnsafeFast(x)
push   esi
call   RtlReleaseActivationContext(x)
retn
; END OF FUNCTION CHUNK FOR _RtlDispatchAPC@12
```

(./AtomBombing\_Brand New Code Injection for Windows - Breaking Malware\_files/Figure-3-RtlDispatchApc-3.png)

Figure 4: ntdll!RtlDispatchAPC – RtlDeactivateActivationContextUnsafeFast

If, on the other hand, no activation context needs to be activated:



```
push [ebp+p_ApcProcParam]
mov  ecx, [ebp+p_ApcProc]
call ds:___guard_check_icall_fptr
call [ebp+p_ApcProc]
```

(./AtomBombing\_Brand New Code Injection for Windows - Breaking Malware\_files/Figure-4-RtlDispatchApc-4.png)

Figure 5: ntdll!RtlDispatchAPC – no activation context

The code skips all the activation context related stuff and simply dispatches the APC right away after calling CFG.

What does all this mean? When calling QueueUserApc we are forced to pass an APCProc which expects one parameter. However, under the hood QueueUserApc uses NtQueueApcThread to call ntdll!RtlDispatchAPC which expects 3 parameters.

What was our goal? To call GlobalGetAtomName. How many parameters does it expect? 3. Can we do this? Yes. How? NtQueueApcThread!

See [main\\_ApcWriteProcessMemory](https://github.com/BreakingMalwareResearch/atom-bombing/blob/f5a04e3fd6eda114e1a7e66429d7e9ec864ec7d4/AtomBombing/main.cpp#L822) (<https://github.com/BreakingMalwareResearch/atom-bombing/blob/f5a04e3fd6eda114e1a7e66429d7e9ec864ec7d4/AtomBombing/main.cpp#L822>) in [AtomBombing's GitHub repository](https://github.com/BreakingMalwareResearch/atom-bombing) (<https://github.com/BreakingMalwareResearch/atom-bombing>).

## AtomBombing Stage 2: Execution

Obviously I could never hope to consistently find RWX code caves in my target processes. I needed a way to consistently allocate RWX memory in the target process without calling VirtualAllocEx within the context of the injecting process. Sadly, I could not find any such function that I could invoke via APC and would allow me to allocate executable memory or change the protection flags of already allocated memory.

What do we have so far? Write-what-where + a burning desire to get some executable memory. I thought long and hard how to get over this hurdle, and then it hit me. When DEP was invented, its creators thought, "that's it, data is no longer executable, therefore no one will ever be able to exploit vulnerabilities again". Unfortunately, that was not the case; a new exploitation technique was invented solely to bypass DEP: ROP – Return Oriented Programming.

How can we use ROP to our advantage in order to execute our shellcode in the target process?

We can copy our code to an RW code cave in the target process (using the method described in stage 1). Then use a meticulously crafted ROP chain to allocate RWX memory, copy the code from the RW code cave to the newly allocated RWX memory, and finally jump to the RWX memory and execute it.

Finding an RW code cave is not a big problem. For this proof of concept, I decided to use the unused space after the data section of kernelbase.

See [main\\_GetCodeCaveAddress](https://github.com/BreakingMalwareResearch/atom-bombing/blob/f5a04e3fd6eda114e1a7e66429d7e9ec864ec7d4/AtomBombing/main.cpp#L1098) (<https://github.com/BreakingMalwareResearch/atom-bombing/blob/f5a04e3fd6eda114e1a7e66429d7e9ec864ec7d4/AtomBombing/main.cpp#L1098>) in [AtomBombing's GitHub repository](https://github.com/BreakingMalwareResearch/atom-bombing) (<https://github.com/BreakingMalwareResearch/atom-bombing>).

## The ROP Chain:

Our ROP chain needs to do 3 things:

1. Allocate RWX memory
2. Copy the shellcode from the RW code cave to the newly allocated RWX memory
3. Execute the newly allocated RWX memory

## ROP Chain Step 1: Allocating RWX Memory

We would like to allocate some RWX memory. The first function that comes to mind is `VirtualAlloc` – a very useful function that can be used to allocate RWX memory. The only problem is that the function returns the newly allocated RWX memory in `EAX` which would make our ROP chain complicated by having to find a way to pass the value `VirtualAlloc` stored in `EAX` to the next function in the chain.

A very neat trick can be employed in order to simplify our ROP chain and make it more sophisticated. Instead of using `VirtualAlloc`, we can use `ZwAllocateVirtualMemory`, which returns the newly allocated RWX memory as an output parameter. This way we can actually set up our stack so that `ZwAllocateVirtualMemory` stores the newly allocated memory further along the stack, effectively passing the address to the next function in the chain (see Table 1).

## ROP Chain Step 2: Copying the Shellcode

The next function we need is a function that will copy memory from one buffer to another. Two options come to mind: `memcpy` and `RtlMoveMemory`. When creating this kind of ROP chain one might be initially inclined to go with `RtlMoveMemory` because it uses the `stdcall` calling convention, meaning it will clean up the stack after itself. This is a special case though. We need to copy memory to an address (placed on the stack by `ZwAllocateVirtualMemory`) and then somehow this address needs to be called. If we used `RtlMoveMemory`, it will pop the address of the RWX shellcode right off the stack upon its return. On the other hand, if we use `memcpy`, the first entry on the stack would be the return address of `memcpy`, followed by the destination parameter of `memcpy` (i.e. the RWX shellcode).

## ROP Chain Step 3: Executing the newly allocated RWX memory

We have allocated RWX memory and copied our shellcode to it. We are about to return from `memcpy` but the address of the RWX shellcode on the stack is 4 bytes away from the return address. Therefore, all we have to do is add an extremely simple gadget to our ROP chain. This simple gadget will execute the opcode “ret”. `memcpy` will return to this simple gadget which will “ret” right into our RWX shellcode.

See [main\\_FindRetGadget \(https://github.com/BreakingMalwareResearch/atom-bombing/blob/f5a04e3fd6eda114e1a7e66429d7e9ec864ec7d4/AtomBombing/main.cpp#L1131\)](https://github.com/BreakingMalwareResearch/atom-bombing/blob/f5a04e3fd6eda114e1a7e66429d7e9ec864ec7d4/AtomBombing/main.cpp#L1131) in [AtomBombing's GitHub repository \(https://github.com/BreakingMalwareResearch/atom-bombing\)](https://github.com/BreakingMalwareResearch/atom-bombing).

For those who have to see it to believe it:

Set `EIP` to point to `ZwAllocateVirtualMemory`, and `ESP` to point to this ROP chain:

Address	Value	Comment
0x30000000	<code>ntdll!memcpy</code>	// Return address from <code>ZwAllocateVirtualMemory</code>
0x30000004	<code>0xffffffff</code>	// Pseudo handle to the current process
0x30000008	<code>0x30000020</code>	// Where to store the allocated memory
0x3000000C	<code>NULL</code>	// Irrelevant
0x30000010	<code>0x30000028</code>	// Pointer to the size of the needed memory
0x30000014	<code>MEM_COMMIT</code>	// Commit and not reserve
0x30000018	<code>PAGE_EXECUTE_READWRITE</code>	// RWX
0x3000001C	<code>POINTER_TO_SOME_RET_INSTRUCTION</code>	// Return Address from <code>memcpy</code> , our extremely simple ret gadget.
0x30000020	<code>NULL</code>	// Where the allocated memory will be saved and the destination parameter of <code>memcpy</code> . This will store the address of the RWX shellcode.

0x30000024	CODE_CAVE_ADDRESS	// The RW code cave containing the shellcode to be copied
0x30000028	SHELLCODE_SIZE	// The size of the shellcode to be allocated

**Table 1: The whole ROP chain.**

See [main BuildROPChain \(https://github.com/BreakingMalwareResearch/atom-bombing/blob/f5a04e3fd6eda114e1a7e66429d7e9ec864ec7d4/AtomBombing/main.cpp#L1192\)](https://github.com/BreakingMalwareResearch/atom-bombing/blob/f5a04e3fd6eda114e1a7e66429d7e9ec864ec7d4/AtomBombing/main.cpp#L1192) in AtomBombing's GitHub repository (<https://github.com/BreakingMalwareResearch/atom-bombing>).

## Invoking the ROP Chain

But wait, APCs allow me to send 3 parameters. Obviously I need to store 11 parameters on the stack. Our best bet is to pivot the stack to some RW memory which will contain our ROP chain (e.g. the RW code cave in kernelbase).

How could we pivot the stack?

```
NTSYSAPI NTSTATUS NTAPI NtSetContextThread(
_In_ HANDLE hThread,
_In_ const CONTEXT *lpContext
);
```

This syscall will set the context (register values) of hThread to the values contained in lpContext. If we can get the target process to call this syscall with an lpContext that will set ESP to point to our ROP chain and set EIP to point to ZwAllocateVirtualMemory, then our ROP chain will execute. The execution of the ROP chain will eventually lead to the execution of our shellcode.

How do we get the target process to make this call? APC has been good to us so far, but this syscall expects 2 parameters and not 3, so when it returns the stack will be corrupt, and the behavior will be undefined. That said, if we pass a handle to the current thread as hThread, then the function will never return. The reason is that once execution gets passed on to the kernel, the context of the thread will be set to the context specified by lpContext, and there will be no trace that NtSetContextThread was ever called. If everything works out as we hope, we will have successfully hijacked a thread and got it to execute our malicious shellcode.

See [main ApcSetThreadContext \(https://github.com/BreakingMalwareResearch/atom-bombing/blob/f5a04e3fd6eda114e1a7e66429d7e9ec864ec7d4/AtomBombing/main.cpp#L877\)](https://github.com/BreakingMalwareResearch/atom-bombing/blob/f5a04e3fd6eda114e1a7e66429d7e9ec864ec7d4/AtomBombing/main.cpp#L877) in AtomBombing's GitHub repository (<https://github.com/BreakingMalwareResearch/atom-bombing>).

## AtomBombing Stage 3: Restoration

We do have one problem, though. The thread that we hijacked had a purpose before we had hijacked it. If we don't restore its execution, there is no telling what kind of effect we could have on the target process.

How do we restore execution? I'd like to remind you that we are now in the context of an APC. When the APC function completes, somehow execution is restored safely. Let's look at the dispatching of APCs from the target process's point of view.

It looks like the function in charge of dispatching APCs (WaitForSingleObjectEx in this example) is ntdll!KiUserApcDispatcher.

```

; Exported entry 109. KiUserApcDispatcher

; Attributes: noreturn

public __stdcall KiUserApcDispatcher(x, x, x, x)
__stdcall KiUserApcDispatcher(x, x, x, x) proc near

Context= CONTEXT ptr 10h
arg_2D8= byte ptr 2DCh

lea     eax, [esp+arg_2D8]
mov     ecx, large fs:0
mov     edx, offset KiUserApcExceptionHandler(x,x,x,x)
mov     [eax], ecx
mov     [eax+4], edx
mov     large fs:0, eax
pop     eax
lea     edi, [esp-4+Context]
mov     ecx, eax
call    ds:__guard_check_icall_fptr
call    ecx
mov     ecx, [edi+2CCh]
mov     large fs:0, ecx
push    1 ; TestAlert
push    edi ; Context
call    ZwContinue(x,x)
mov     esi, eax

```

([./AtomBombing - Brand New Code Injection for Windows - Breaking Malware files/Figure-5-KiUserApcDispatcher.png](#))

Figure 6: KiUserApcDispatcher


We can see 3 “calls” in this block of code. The first call is to CFG, the next call is to ECX (which is the address of the APC function), and finally a call to the undocumented ZwContinue.

ZwContinue expects to receive a pointer to a CONTEXT structure and resumes the execution. Actually the kernel will check if there are any more APCs in the thread’s APC queue, and dispatch them before finally resuming the thread’s original execution, but we can ignore that.

The CONTEXT structure being passed to ZwContinue is stored in EDI before calling the APC function (stored in ECX). We can save EDI’s value at the beginning of our shellcode, and call ZwContinue with EDI’s original value at the end of the shellcode, thereby restoring execution safely.

See *AtomBombingShellcode* (<https://github.com/BreakingMalwareResearch/atom-bombing/blob/f5a04e3fd6eda114e1a7e66429d7e9ec864ec7d4/AtomBombingShellcode/main.c#L31>) in *AtomBombing’s* GitHub repository (<https://github.com/BreakingMalwareResearch/atom-bombing>).

We have to make sure that the value of EDI will not be overridden during the call to NtSetContextThread, since it modifies the values of the registers. This can easily be accomplished by setting ContextFlags (member of the CONTEXT structure passed to NtSetContextThread) to CONTEXT\_CONTROL which means that only EBP, EIP, SEGCS, EFLAGS, ESP, and SEGSS will be affected. As long as (CONTEXT.ContextFlags&CONTEXT\_INTEGER == 0) we should be ok.

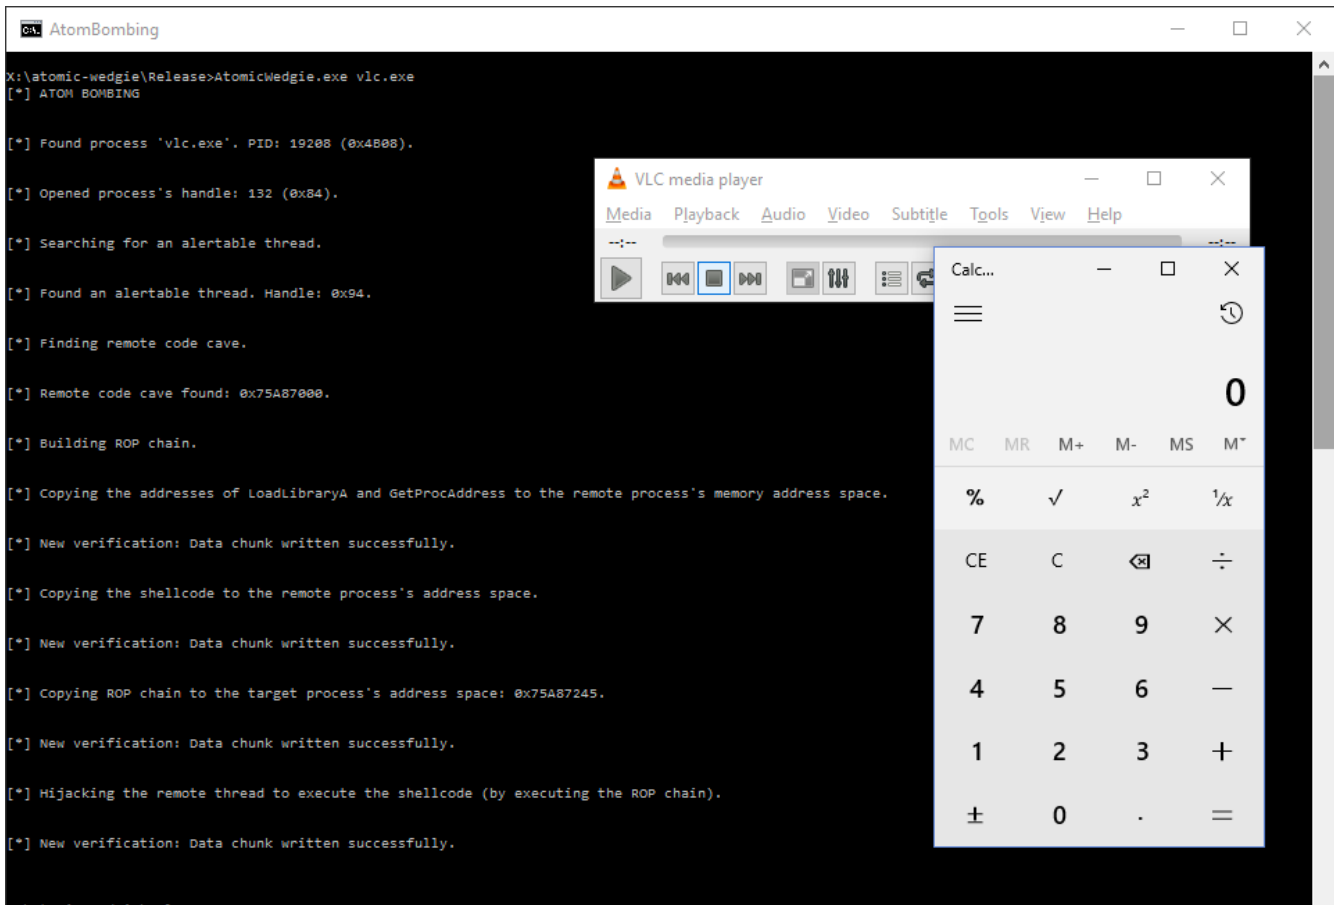


([./AtomBombing - Brand New Code Injection for Windows - Breaking Malware files/Figure-6-Chrome.png](#))

Figure 7: AtomBombing chrome.exe

And there you have it, we have injected code into chrome.exe. Our injected code spawned the classic calc.exe proving that it works.

Let’s try to inject code into vlc.exe:



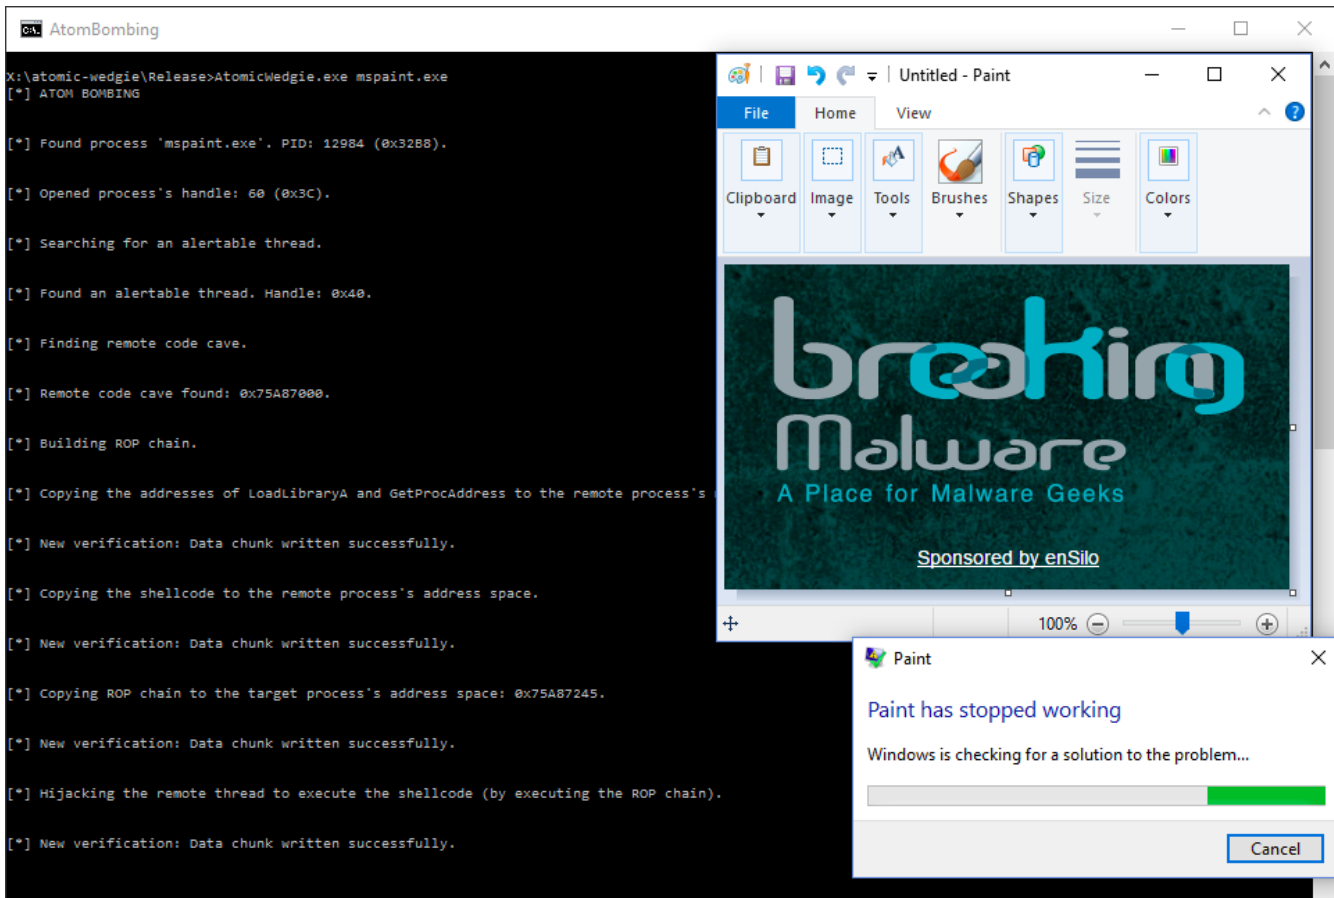
[./AtomBombing - Brand New Code Injection for Windows - Breaking Malware files/Figure-7-VLC.png](#)

Figure 8: AtomBombing vlc.exe

The complete implementation can be found on [GitHub \(https://github.com/BreakingMalwareResearch/atom-bombing\)](https://github.com/BreakingMalwareResearch/atom-bombing). It has been tested against Windows 10 x64 Build 1511 (WOW) and Windows 10 x86 Build 10240. Compile for "release".

Let's do the same with mspaint.exe:





[./AtomBombing\\_Brand New Code Injection for Windows - Breaking Malware\\_files/figure-8-MSPaint.png](#)

Figure 9: AtomBombing mspaint.exe

Oh no, it crashed.

## Final Steps

How do we proceed from here? I have worked it out and at this point, I'd rather leave this as an exercise to the reader. As an initial hint, I suggest you take a look at my previous blog post (<https://breakingmalware.com/documentation/documenting-undocumented-adding-control-flow-guard-exceptions/> (<https://breakingmalware.com/documentation/documenting-undocumented-adding-control-flow-guard-exceptions/>)). I'm sure you'll also find creative ideas, that I myself haven't found, to handle this problem and I'd be happy to start this discussion.

You can use the comments below or catch me [@tal\\_liberman](https://twitter.com/Tal_Liberman) ([https://twitter.com/Tal\\_Liberman](https://twitter.com/Tal_Liberman)). Through Twitter, I'll also release some tidbits throughout the week. At any rate, I will publish my solution next week.

## APPENDIX: Finding Alertable Threads

One thing we have not yet mentioned is that `QueueUserApc` only works on threads that are in an alertable state. How does a thread enter an alertable state?

According to Microsoft:

""

A thread can only do this by calling one of the following functions with the appropriate flags:

- `SleepEx`
- `WaitForSingleObjectEx`
- `WaitForMultipleObjectsEx`
- `SignalObjectAndWait`
- `MsgWaitForMultipleObjectsEx`

When the thread enters an alertable state, the following events occur:

1. The kernel checks the thread's APC queue. If the queue contains callback function pointers, the kernel removes the pointer from the queue and sends it to the thread.
2. The thread executes the callback function.
3. Steps 1 and 2 are repeated for each pointer remaining in the queue.
4. When the queue is empty, the thread returns from the function that placed it in an alertable state.

[https://msdn.microsoft.com/en-us/library/windows/desktop/aa363772\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/desktop/aa363772(v=vs.85).aspx) ([https://msdn.microsoft.com/en-us/library/windows/desktop/aa363772\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/desktop/aa363772(v=vs.85).aspx))

""

For our technique to be effective the target process must have at least one thread that is in an alertable state, or will enter an alertable state at some point, otherwise our APCs will never actually execute.

I've checked various software, and I've noticed that most of the programs I've checked have at least one alertable thread. Examples: Chrome.exe, Iexplore.exe, Skype.exe, VLC.exe, MsPaint.exe, WmiPrvSE.exe, etc.

So now we have to find an alertable thread in the target process. There are many ways of doing this. I chose to use a method that is trivial, works in most cases, and is easy to implement and explain.

We'll create an event for each thread in the target process, then ask each thread to set its corresponding event. We'll wait on the event handles, until one is triggered. The thread whose corresponding event was triggered is an alertable thread.

How can an event be set? By calling `SetEvent(HANDLE hEvent)`.

How will we get the threads in the target process to call `SetEvent`? APC of course. Since `SetEvent` receives exactly one parameter we can use `QueueUserApc` to call it. The actual details of the implementation can be found in [main FindAlertableThread \(https://github.com/BreakingMalwareResearch/atom-bombing/blob/f5a04e3fd6eda114e1a7e66429d7e9ec864ec7d4/AtomBombing/main.cpp#L1494\)](https://github.com/BreakingMalwareResearch/atom-bombing/blob/f5a04e3fd6eda114e1a7e66429d7e9ec864ec7d4/AtomBombing/main.cpp#L1494) in [AtomBombing's GitHub repository \(https://github.com/BreakingMalwareResearch/atom-bombing\)](https://github.com/BreakingMalwareResearch/atom-bombing).

[Injection Techniques \(https://breakingmalware.com/category/injection-techniques/\)](https://breakingmalware.com/category/injection-techniques/)

[apc \(https://breakingmalware.com/tag/apc/\)](https://breakingmalware.com/tag/apc/), [atombombing \(https://breakingmalware.com/tag/atombombing/\)](https://breakingmalware.com/tag/atombombing/), [code injection \(https://breakingmalware.com/tag/code-injection/\)](https://breakingmalware.com/tag/code-injection/), [research \(https://breakingmalware.com/tag/research/\)](https://breakingmalware.com/tag/research/), [windows \(https://breakingmalware.com/tag/windows/\)](https://breakingmalware.com/tag/windows/)

## Post navigation


[← Elastic Boundaries – Elevating privileges by environment variables expansion \(https://breakingmalware.com/vulnerabilities/elastic-boundaries-elevating-privileges-by-environment-variables-expansion/\)](https://breakingmalware.com/vulnerabilities/elastic-boundaries-elevating-privileges-by-environment-variables-expansion/)

[AtomBombing CFG Protected Processes → \(https://breakingmalware.com/injection-techniques/atombombing-cfg-protected-processes/\)](https://breakingmalware.com/injection-techniques/atombombing-cfg-protected-processes/)

We were unable to load Disqus. If you are a moderator please see our [troubleshooting guide](#).

Search for:

## Tomer Bitton, VP Research, enSilo



Obsessed by malware, hostile code and extreme packers. Low-level researcher. Feel free to contact.



<http://il.linkedin.com/in/p4r4n0id>



[https://twitter.com/p4r4n0id\\_il](https://twitter.com/p4r4n0id_il)