

# Introduction to Intel x86-64 Assembly, Architecture, Applications, & Alliteration

Xeno Kovah – 2014-2015  
[xeno@legbacore.com](mailto:xeno@legbacore.com)

# All materials is licensed under a Creative Commons “Share Alike” license.

- <http://creativecommons.org/licenses/by-sa/3.0/>

## You are free:



to **Share** — to copy, distribute and transmit the work



to **Remix** — to adapt the work

## Under the following conditions:



**Attribution** — You must attribute the work in the manner specified by the author or licensor (but not in any way that suggests that they endorse you or your use of the work).



**Share Alike** — If you alter, transform, or build upon this work, you may distribute the resulting work only under the same, similar or a compatible license.

Attribution condition: You must indicate that derivative work  
"Is derived from Xeno Kovah's 'Intro x86-64' class, available at <http://OpenSecurityTraining.info/IntroX86-64.html>"

Attribution condition: You must indicate that derivative work

"Is derived from Xeno Kovah's 'Intro x86-64' class, available at <http://OpenSecurityTraining.info/IntroX86-64.html>"

# ArrayLocalVariable2.c

Zero-initializing the array

|                          |                   |          |                           |
|--------------------------|-------------------|----------|---------------------------|
|                          | main:             |          |                           |
|                          | 00000000140001000 | push     | rdi                       |
|                          | 00000000140001002 | sub      | rsp,20h                   |
|                          | 00000000140001006 | xor      | eax, eax                  |
| //ArrayLocalVariable2.c: | 00000000140001008 | mov      | word ptr [rsp+8], ax      |
| short main({             | 0000000014000100D | lea      | rax, [rsp+0Ah]            |
| int a;                   | 00000000140001012 | mov      | rdi, rax                  |
| short b[6] = {0};        | 00000000140001015 | xor      | eax, eax                  |
| a = 0x100d;              | 00000000140001017 | mov      | ecx, 0Ah                  |
| b[1] = (short)a;         | 0000000014000101C | rep stos | byte ptr [rdi]            |
| return b[1];             | 0000000014000101E | mov      | dword ptr [rsp], 100Dh    |
| }                        | 00000000140001025 | mov      | eax, 2                    |
|                          | 0000000014000102A | imul     | rax, rax, 1               |
|                          | 0000000014000102E | movzx    | ecx, word ptr [rsp]       |
|                          | 00000000140001032 | mov      | word ptr [rsp+rax+8], cx  |
|                          | 00000000140001037 | mov      | eax, 2                    |
|                          | 0000000014000103C | imul     | rax, rax, 1               |
|                          | 00000000140001040 | movzx    | eax, word ptr [rsp+rax+8] |
|                          | 00000000140001045 | add      | rsp, 20h                  |
|                          | 00000000140001049 | pop      | rdi                       |
|                          | 0000000014000104A | ret      |                           |



# REP STOS - Repeat Store String

- STOS is one of number of instructions that can have the “rep” prefix added to it, which repeat a single instruction multiple times.
- All rep operations use `*cx` register as a “counter” to determine how many times to loop through the instruction. Each time it executes, it decrements `*cx`. Once `*cx == 0`, it continues to the next instruction.
- Either stores 1, 2, 4, or 8 bytes at a time
- Either fill 1 byte at `[di]` with `al` or fill 1/2/4/8 bytes at `[*di]` with `al/*ax`
- Moves the `*di` register forward 1/2/4/8 bytes at a time, so that the repeated store operation is storing into consecutive locations.
- So there are 3 pieces which must happen before the actual rep stos occurs: set `*di` to the start destination, `*ax/al` to the value to store, and `*cx` to the number of times to store

As with other instructions prefixes like “LOCK”, “REP” can only be used with certain instructions - as defined in the manual.

# ArrayLocalVariable2.c takeaways

- If you're manually coding asm, REP STOS is functionally a memset()
- Sometimes when you use memset() from C, the compiler may turn it into a REP STOS

//ArrayLocalVariable2.c:

```
short main(){
    int a;
    short b[6] = {0};
    a = 0x100d;
    b[1] = (short)a;
    return b[1];
}
```

```
main:
    push    rdi
    sub     rsp,20h
    xor     eax,eax
    mov     word ptr [rsp+8],ax
    lea     rax,[rsp+0Ah]
    mov     rdi,rax
    xor     eax,eax
    mov     ecx,0Ah
    rep stos byte ptr [rdi]
    mov     dword ptr [rsp],100Dh
    mov     eax,2
    imul    rax,rax,1
    movzx   ecx,word ptr [rsp]
    mov     word ptr [rsp+rax+8],cx
    mov     eax,2
    imul    rax,rax,1
    movzx   eax,word ptr [rsp+rax+8]
    add     rsp,20h
    pop     rdi
    ret
```

# ThereWillBe0xb100d.c

```
int main(){
    char buf[40];
    buf[39] = 42;
    return 0xb100d;
}
```

# ThereWillBe0xb100d.c

main:

```
00000000140001010 push    rdi
00000000140001012 sub     rsp,60h
00000000140001016 mov     rdi,rsp
00000000140001019 mov     ecx,18h
0000000014000101E mov     eax,0CCCCCCCCh
00000000140001023 rep stos dword ptr [rdi]
00000000140001025 mov     eax,1
0000000014000102A imul    rax,rax,27h
0000000014000102E mov     byte ptr [rsp+rax+28h],2Ah
00000000140001033 mov     eax,0xb100d
00000000140001038 mov     edi,eax
0000000014000103A mov     rcx,rsp
0000000014000103D lea     rdx,[__xi_z+1A0h (0140006910h)]
00000000140001044 call    _RTC_CheckStackVars (01400010B0h)
00000000140001049 mov     eax,edi
0000000014000104B add     rsp,60h
0000000014000104F pop     rdi
00000000140001050 ret
```

# rep stos setup

```
00000000140001016 mov    rdi, rsp
```

**Set rdi - the destination**

```
00000000140001019 mov    ecx, 18h
```

**Set ecx - the count**

```
0000000014000101E mov    eax, 0CCCCCCCCh
```

**Set eax - the value**

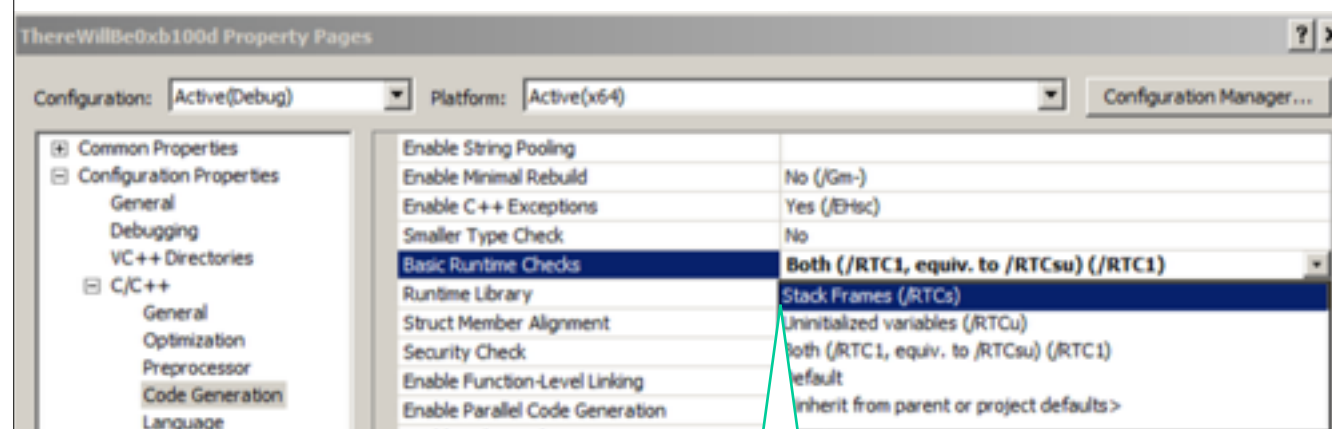
```
00000000140001023 rep stos dword ptr [rdi]
```

**Start the repeated store**

- So what's this going to do? Store 0x18 copies of the dword 0xCCCCCCCC starting at rsp
- And that just happens to be 0x60 bytes of 0xCC, the entire reserved stack space!



Q: Where does the rep stos come from in this example?



A: Compiler-auto-generated code. From the stack frames runtime check option. This is enabled by default in the debug build. Disabling this option removes the compiler-generated code.

# More straightforward without the runtime check

main:

```
00000000140001010 sub    rsp,38h
00000000140001014 mov     eax,1
00000000140001019 imul    rax,rax,27h
0000000014000101D mov     byte ptr [rsp+rax],2Ah
00000000140001021 mov     eax,0B100Dh
00000000140001026 add     rsp,38h
0000000014000102A ret
```

But still not entirely clear :)

# Instructions we now know (29)

- NOP
- PUSH/POP
- CALL/RET
- MOV
- ADD/SUB
- IMUL
- MOVZX/MOVSX
- LEA
- JMP/Jcc (family)
- CMP/TEST
- AND/OR/XOR/NOT
- INC/DEC
- SHR/SHL/SAR/SAL
- DIV/IDIV
- REP STOS