

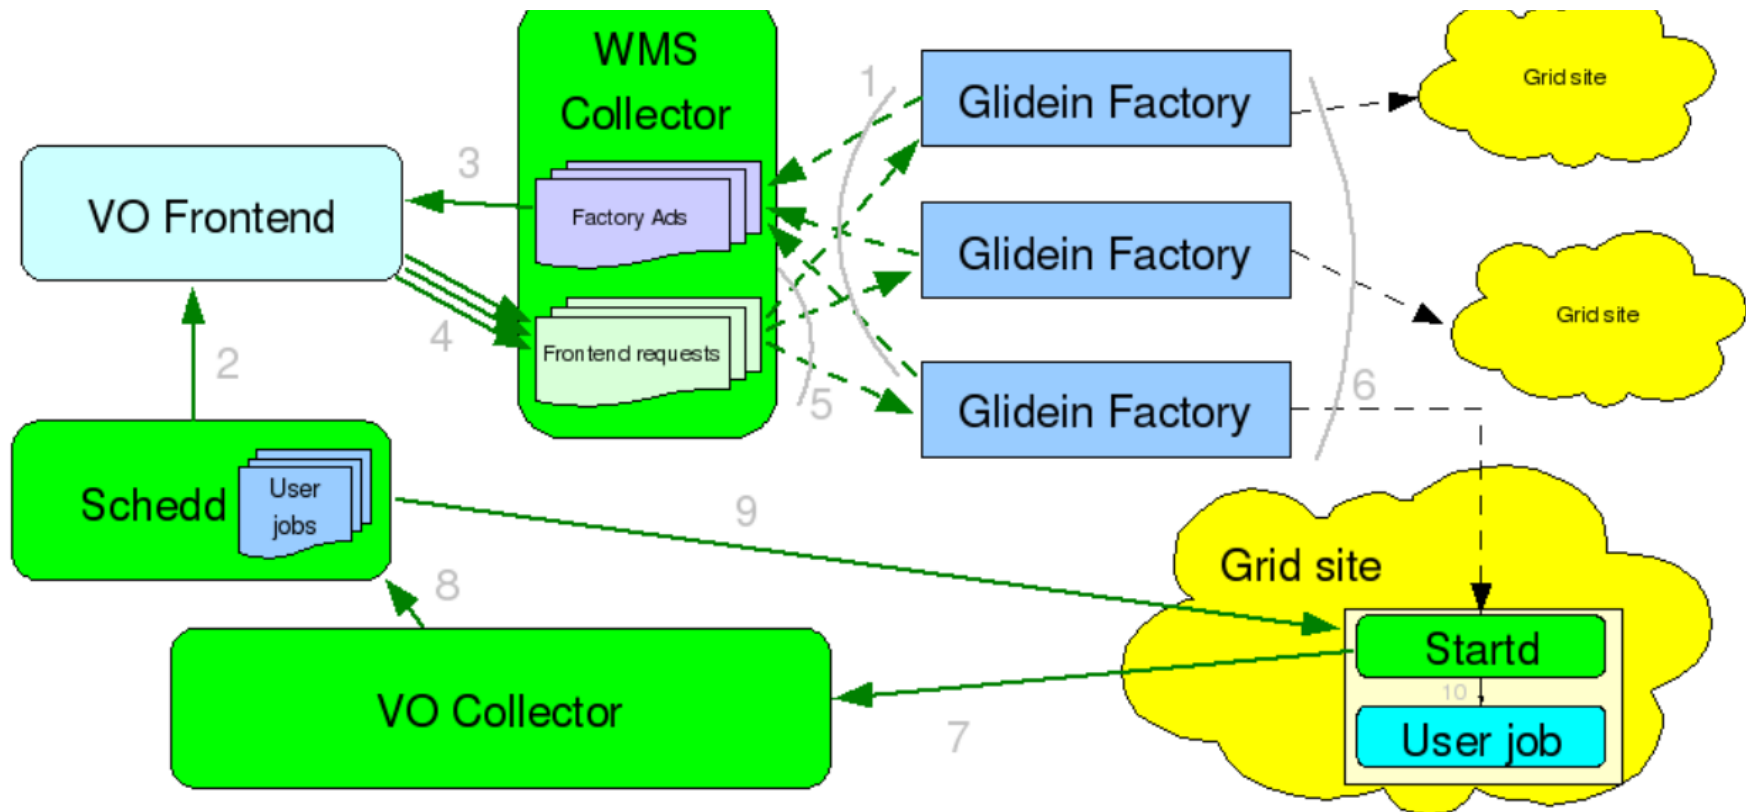


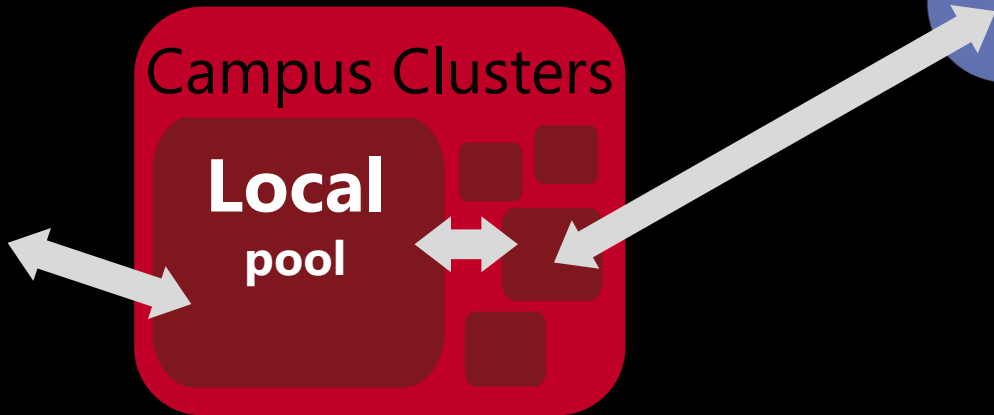


MicroArch

Nomen Numen

Greg Thain

OS Pool












The
beginning of
Wisdom
is to call things by their proper
name.

-- Confucius

NATIONAL BESTSELLER

THE
NAME
OF THE
ROSE

UMBERTO
ECO



PATRICK
ROTHFUSS

THE
NAME
OF THE
WIND

The Power of Clarity


What we say

Most of the jobs work, but some are failing in the job router, but those jobs that do work are running faster than we expected.




What they hear

Blah blah blah JOBS blah
blah blah JOBS blah blah
blah JOBS




The Power of Metaphor


Tree representation



Array representation



Heap




AGENDA	
MEETING NAME:	
DATE:	TIME:
LOCATION:	
ATTENDEES:	
NEXT MEETING:	
TOPIC	FEEDBACK

The Power of Generation




The Best Time to Plant a Tree

Submit
Point


















Access
Point
(AP)

Compute
Element
(CE)




Compute
Entryway
(CE)





Single-Threaded Floating-Point Performance


Based on adjusted SPECfp® results



Year	Insn	Family	Description
<u>2001</u>	<u>SSE2</u>	<u>P6</u>	Attempt to replace the original <u>MMX</u> instructions; wider <u>XMM</u> registers offer better performance
<u>2004</u>	<u>SSE3</u>	<u>Core</u>	Streaming SIMD Extensions 3; new instructions to operate on multiple values in the same register
<u>2006</u>	<u>SSSE3</u>	<u>Core</u>	Supplemental Streaming SIMD Extensions 3; additional instructions for working with packed integers
<u>2007</u>	<u>SSE4.1</u>	<u>Penryn</u>	Streaming SIMD Extensions 4.1; initial set of instructions for manipulating media data
<u>2007</u>	<u>SSE4a</u>	<u>K10</u>	4 SIMD instructions, not related to 4.1/4.2
<u>2007</u>	<u>ABM</u>	<u>K10</u>	Advanced Bit Manipulation; bit counting instructions
<u>2008</u>	<u>SSE4.2</u>	<u>Nehalem</u>	Streaming SIMD Extensions 4.2; second set of instructions for manipulating media data
<u>2007</u>	<u>SSE5</u>		Proposed by AMD in 2007 but was never implemented
<u>2008</u>	<u>SSE4</u>		Streaming SIMD Extensions 4; Another name for SSE4.1 + SSE4.2

<u>2013</u>	<u>FMA3</u>	<u>Haswell</u>	3-operands fused multiply-add
<u>2013</u>	<u>TSX</u>	<u>Haswell</u>	Transactional Synchronization Extensions; adds transactional memory support
<u>2013</u>	<u>AVX2</u>	<u>Haswell</u>	Advanced Vector Extensions; additional instructions
<u>2014</u>	<u>ADX</u>	<u>Broadwell</u>	Multi-Precision Add-Carry Instruction extension
<u>2014</u>	<u>RdRand</u>	<u>Broadwell</u>	Part of Secure Key Technology extension (RdRand, RDSEED)
<u>2014</u>	<u>PREFETCH</u> <u>CH</u>	<u>Broadwell</u>	PREFETCH instructions (previously part of <u>3DNow!</u>)
<u>2015</u>	<u>AVX-512</u>	<u>Airmont</u>	512 bit register operations
<u>2015</u>	<u>MPX</u>	<u>Skylake</u>	Memory Protection Extensions
<u>2015</u>	<u>SGX</u>	<u>Skylake</u>	Software Guard Extensions
<u>2016</u>	<u>SHA</u>	<u>Goldmont</u>	SHA Extensions
<u>2017</u>	<u>SME</u>	<u>Zen</u>	Secure Memory Extensions
<u>2019</u>	<u>TME</u>	<u>Ice Lake</u>	Total Memory Encryption

The hard way...



X86-64-v1

X86-64-v2

X86-64-v3

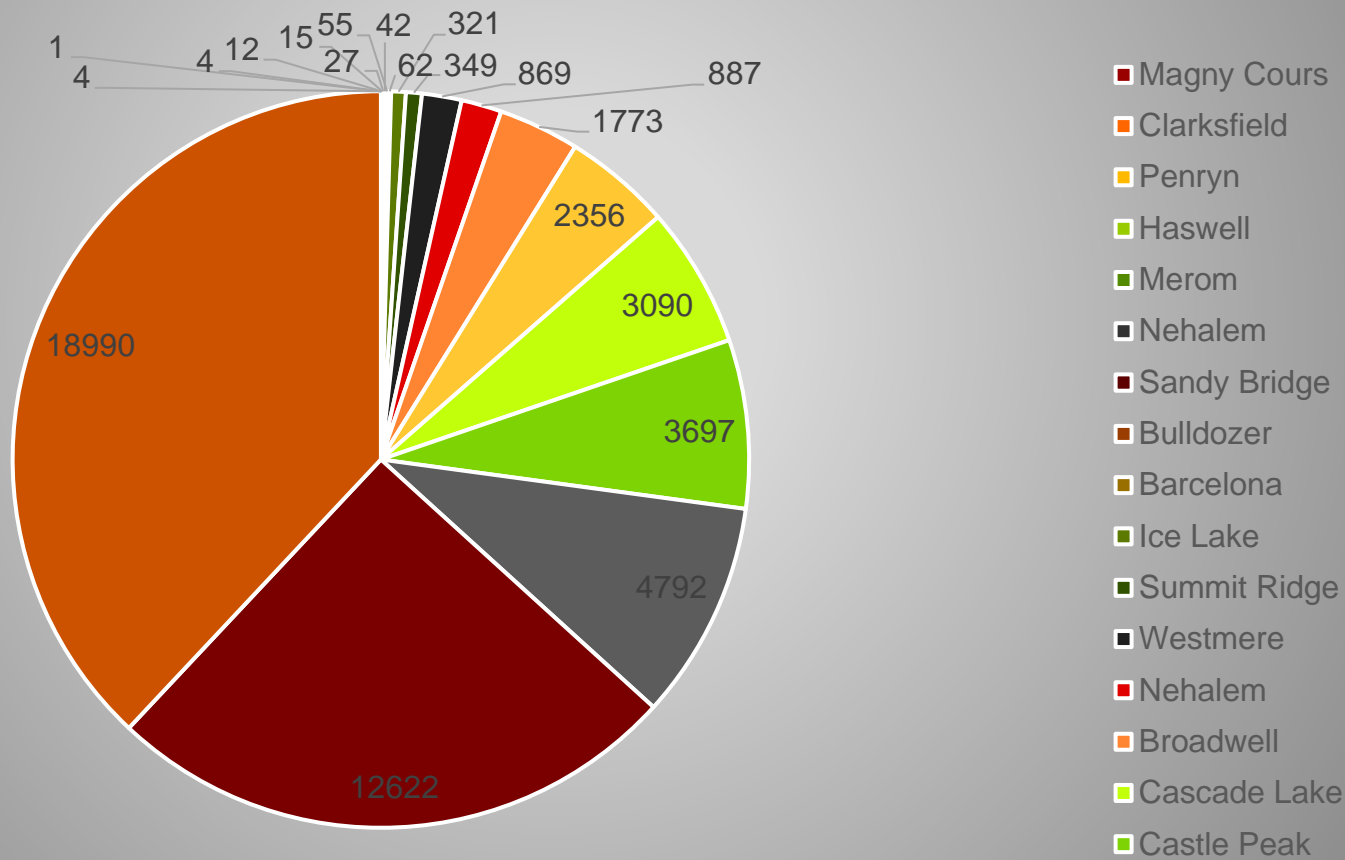
X86-64-v4

OpenSSL

Cryptography and SSL/TLS
Toolkit

Cores in OSPool by codename

Intel/AMD CPU code names in OSPool



Why is this a problem?

- 1) If (CPU job built for) $>$ CPU
→ Crash with SIGILL
- 2) if (CPU job built for) $<$ CPU
→ Run slower than possible

Intel + RedHat to the rescue!

[llvm-dev] New x86-64 micro-architecture levels

Florian Weimer via llvm-dev [llvm-dev at lists.llvm.org](mailto:llvm-dev@lists.llvm.org)

Fri Jul 10 10:30:09 PDT 2020

- Previous message: [\[llvm-dev\] \[LNT\] Build bot problems due to Python dependencies](#)
- Next message: [\[llvm-dev\] New x86-64 micro-architecture levels](#)
- Messages sorted by: [\[date\]](#) [\[thread\]](#) [\[subject\]](#) [\[author\]](#)

Most Linux distributions still compile against the original x86-64 baseline that was based on the AMD K8 (minus the 3DNow! parts, for Intel EM64T compatibility).

There has been an attempt to use the existing AT_PLATFORM-based loading mechanism in the glibc dynamic linker to enable a selection of optimized libraries. But the general selection mechanism in glibc is problematic:

```
hwcaps subdirectory selection in the dynamic loader  
<https://sourceware.org/pipermail/libc-alpha/2020-May/113757.html>
```

We also have the problem that the glibc version of "haswell" is distinct from GCC's -march=haswell (and presumably other compilers):

```
Definition of "haswell" platform is inconsistent with GCC  
<https://sourceware.org/bugzilla/show\_bug.cgi?id=24080>
```

And that the selection criteria are not what people expect:

Microarch: four levels

Haswell	Core	Sapphire Rapids	Nehalem
Broadwell	Core 2	Rapids Lake	Westmere
Skylake	Penryn	Granite Rapids	Sandy Bridge
Cascade Lake	Wolfdale	Willow Cove	Ivy Bridge
...

Microarch: four levels

Core
Core 2
Penryn
Wolfdale
...

x86_64-v1

Nehalem
Westmere
Sandy Bridge
Ivy Bridge
...

x86_64-v2

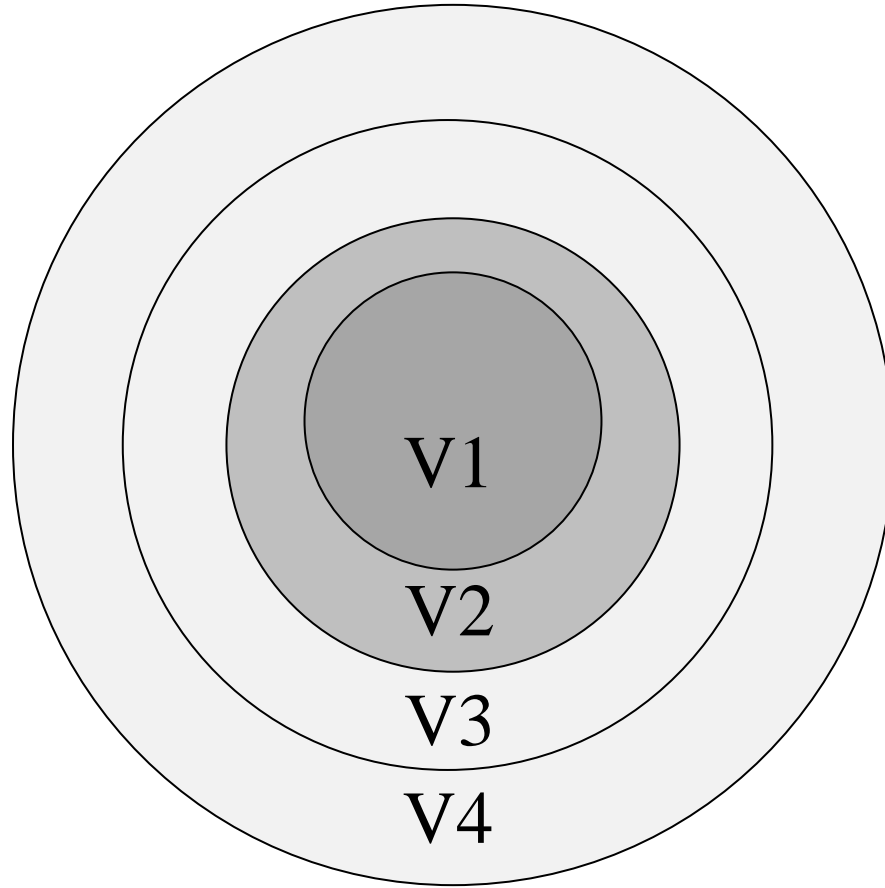
Haswell
Broadwell
Skylake
Cascade Lake
...

x86_64-v3

Sapphire Rapids
Rapids Lake
Granite Rapids
Willow Cove
...

x86_64-v4

Microarch: strict hierarchy



Compilers and microarch

```
$ gcc -march=x86-64-v3 source.cpp
```

```
$ file a.out
```

```
a.out: ELF 64-bit LSB shared object, x86-64, version 1  
(SYSV), dynamically linked, interpreter /lib64/ld-  
linux-x86-64.so.2, for GNU/Linux 3.2.0
```

HTCSS Support for microarch

```
$ condor_status -af Name MicroArch  
slot3@build4000.chtc.wisc.edu x86_64-v3  
slot1@e2007.chtc.wisc.edu x86_64-v4  
slot1@e2008.chtc.wisc.edu x86_64-v4  
slot1@e3001.chtc.wisc.edu x86_64-v1
```


HTCSS Support for microarch

```
$ condor_status -const \  
    'microarch > "x86_64-v2" '  
slot3@build4000.chtc.wisc.edu x86_64-v3  
slot1@e2007.chtc.wisc.edu x86_64-v4  
slot1@e2008.chtc.wisc.edu x86_64-v4
```


HTCSS Support for microarch

```
Executable = calculate.exe  
Requirements = microarch > "x86_64-v2"  
Request_memory = 1Gb  
Request_disk = 100Mb  
Request_cpus = 2  
queue
```

Microarch in OSPool



■ x86_64-v1

■ x86_64-v2

■ x86_64-v3

■ x86_64-v4

More uArch problems

AES acceleration in HTCondor v10.0

Encryption FREE for "modern CPUs"...

Future Wishlist

Automatic detection of microarch in binaries

Microarch specific packaging?



Recap

Names have power

MicroArch

Learn it

Use it

Teach it

Invest in good naming