

Common Optimization Mistakes

DevConf 2010

Ilia Alshanetsky
<http://ilia.ws>

Premature
Optimization

=



Solve the business case,
before optimizing the
solution

Don't Over Engineer

- Understand your audience
- Estimate the scale and growth of your application (based on facts, not marketing fiction)
- Keep timelines in mind when setting the project scope



Simplify, Simplify & Simplify!

- Break complex tasks into simpler sub-components
- Don't be afraid to modularize the code
- More code does not translate to slower code (common misconception)

PHP has grown from less than 1 million LOC to over 2 million LOC since 2000 and has become at least 4 times faster.


Linux kernel code base increase by 40% since 2005 and still managed to improve performance by roughly the same margin.

LOC stats came from ohloh.net

Hardware is Cheaper!



VS




In most cases applications can gain vast performance gains by improving hardware, quickly rather than through slow, error prone code optimization efforts.

Hardware

- CPU bottlenecks can be resolved by more cores and / or CPUs. Typically each year yields a 20-30% speed improvement over past year's CPU speeds.
- Ability to handle large amounts of traffic is often hampered by limited RAM, and thanks to 64bit, each new server can have tons of it.

Hardware

- Drives are often the most common bottleneck, fortunately between RAID and Solid State you can solve that pretty easily now a days.



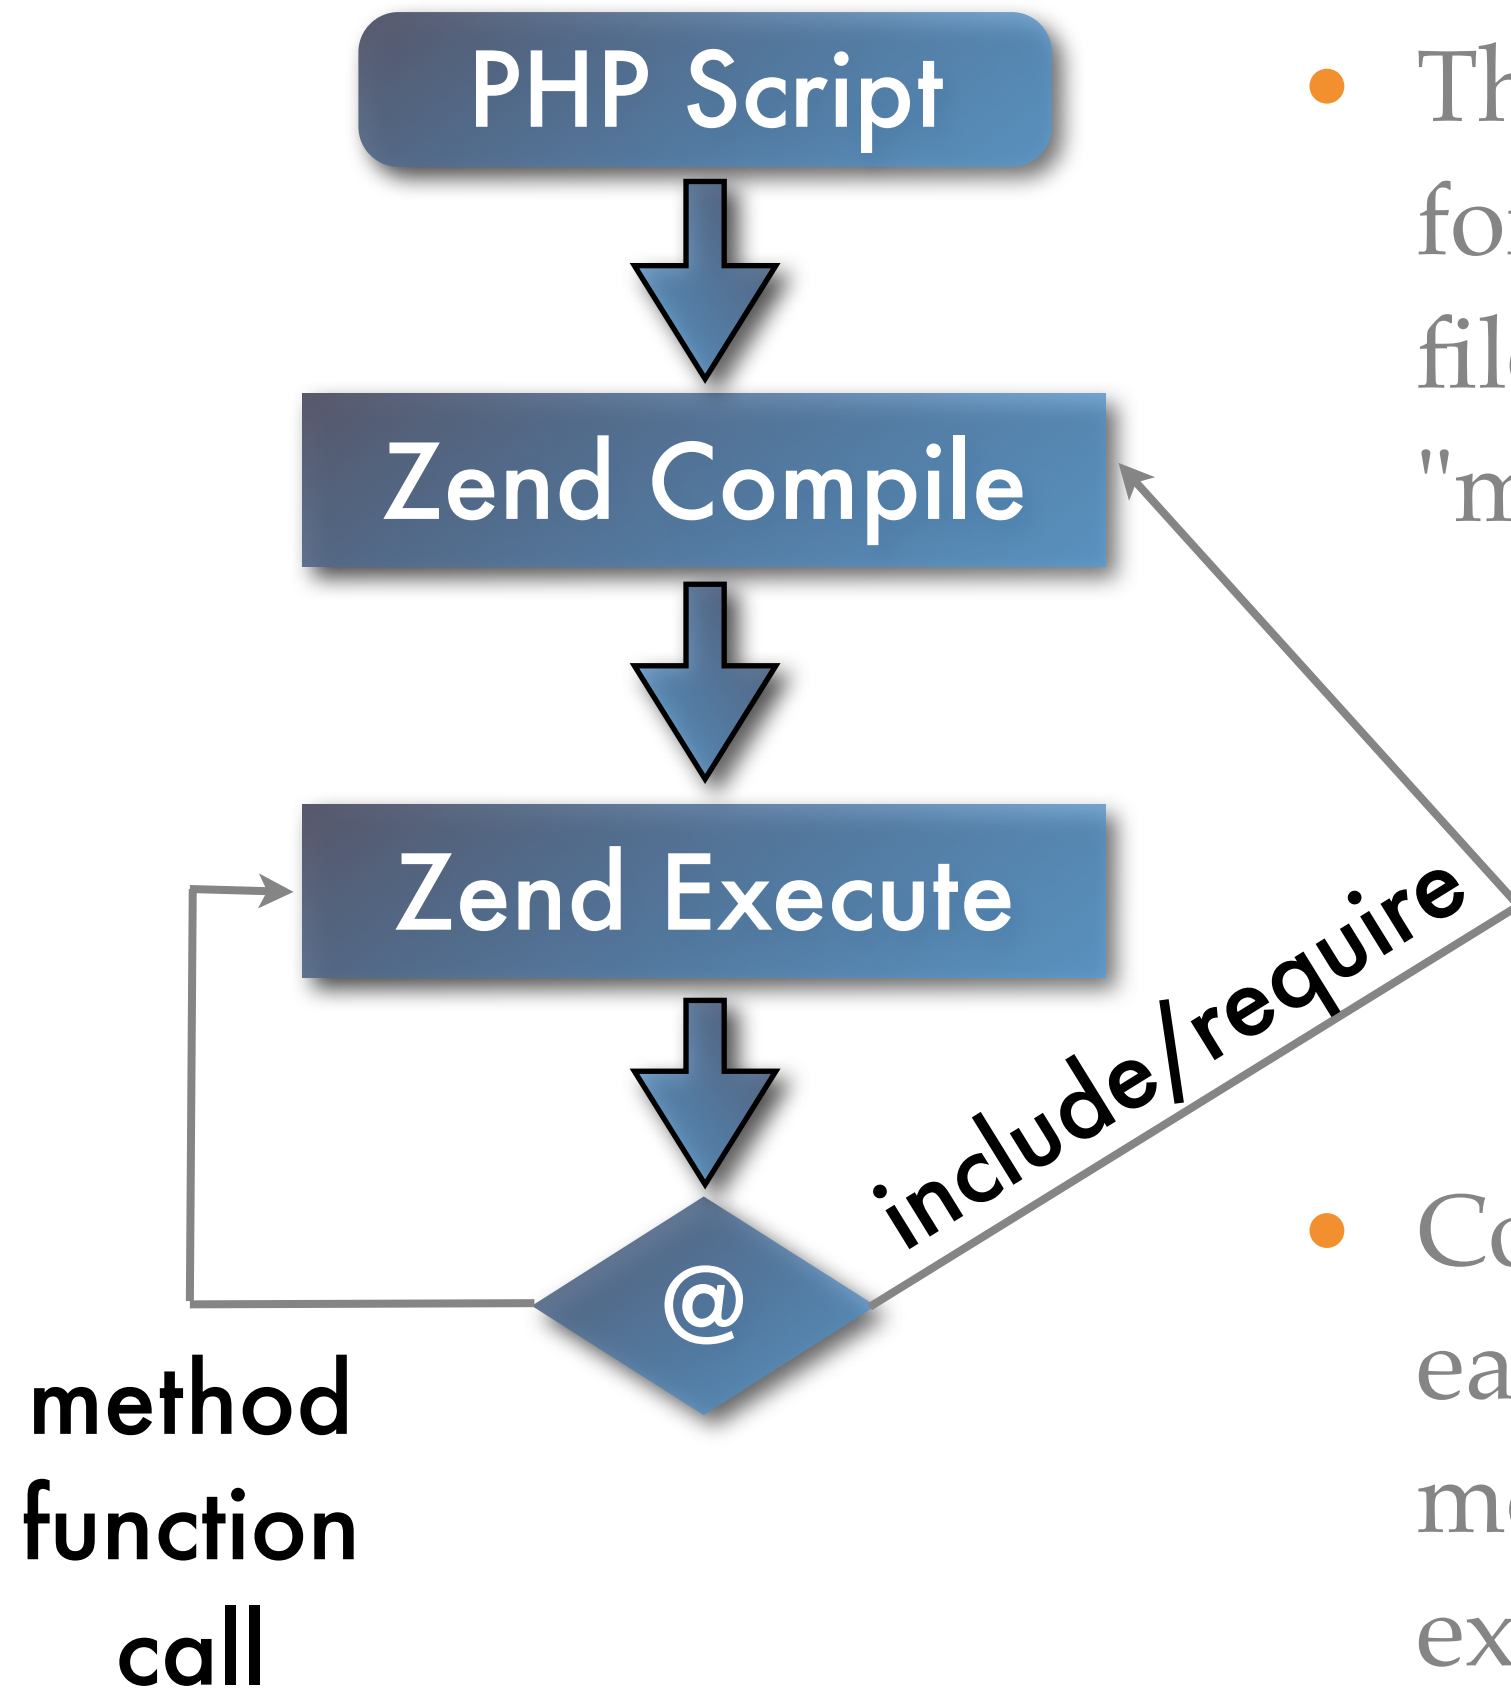
Hardware Caveat

- While quick to give results, in some situations it will not help for long:
 - Database saturation
 - Non-scalable code base
 - Network bound bottleneck
 - Extremely low number sessions per server

Optimize, but don't touch the code

- Typically introduces substantial efficiencies
- Does not endanger code integrity
- Usually simple and quick to deploy
- In the event of problems, often simple to revert

How PHP works in 30 seconds




- This cycle happens for every include file, not just for the "main" script.

- Compilation can easily consume more time than execution.

Opcode Cache

- Each PHP script is interpreted only once for each revision.
- Reduced File IO, opcodes are being read from memory instead of being parsed from disk.
- Opcodes can optimized for faster execution.
- Yields a minimum 20-30% speed improvement and often as much as 200-400%

Quick Comparison



Use In-Memory Caches

- In-memory session storage is MUCH faster than disk or database equivalents.
- Very simple via memcache extension

```
session.save_handler = "memcache"
```

```
session.save_path = "tcp://localhost:11211"
```

Also allows scaling across multiple servers for improved reliability and performance.

Everything has to be Real-time



Complete Page Caching

- Caching Proxy (Ex. Squid)
- Page pre-generation
- On-demand caching

Partial Cache - SQL

- In most applications the primary bottleneck can often be traced to “database work”.
- Caching of SQL can drastically reduce the load caused by unavoidable, complex queries.

SQL Caching Example

```
$key = md5("some sort of sql query");  
if (!($result = memcache_get($key))) {  
    $result = $pdo->query($qry)->fetchAll();  
    // cache query result for 1 hour  
    memcache_set($key, $result, NULL, 3600);  
}
```

Partial Cache - Code

- Rather than optimizing complex PHP operations, it is often better to simply cache their output for a period of time.
 - Faster payoff
 - Lower chance of breaking the code
 - Faster than any “code optimization”

Code Caching Example

```
function myFunction($a, $b, $c) {  
    $key = __FUNCTION__ . serialize(func_get_args());  
    if (!($result = memcache_get($key))) {  
        $result = // function code  
        // cache query result for 1 hour  
        memcache_set($key, $result, NULL, 3600);  
    }  
    return $result;  
}
```


Compile Your Tools

- Distribution binaries suck!
- More often than not you can realize 10-15% speed increase by compiling your own Apache / PHP / Database from source. (*unless you are using Gentoo*)

Output Buffering

- Don't fear output buffering because it uses ram, ram is cheap. IO, not so much.


Matching Your IO Sizes



- The goal is to pass off as much work to the kernel as efficiently as possible.
- Optimizes PHP to OS Communication
- Reduces Number Of System Calls


PHP: Output Control

- Efficient
- Flexible
- In your script, with `ob_start()`
- Everywhere, with `output_buffering = Xkb`
- Improves browser's rendering speed



Apache: Output Control


- The idea is to hand off entire page to the kernel without blocking.



- Set **SendBufferSize** = **PageSize**

OS: Output Control

OS (Linux)



```
/proc/sys/net/ipv4/tcp_wmem
```

```
4096      16384    maxcontentsize
```

```
min        default    max
```


```
/proc/sys/net/ipv4/tcp_mem
```

```
(maxcontentsize * maxclients) / pagesize
```

*** Be careful on low memory systems!**

Upgrade Your PHP

- Before “upgrading” your code, upgrade your PHP. Newer versions are typically faster!



Don't Assume

Assume nothing,
profile everything!




- One of the most common mistakes made even by experienced developers is starting to optimize code without identifying the bottleneck first.

Profile, Profile & Profile

Xdebug and XHProf extensions provide a very helpful mechanism for identifying TRUE bottlenecks in your code.

Kcachegrind



Xdebug provides kcachegrind analyzable output that offers an easy visual overview of your performance problems

Database before code

- One of the most common mistakes people make is optimizing code before even looking at the database.
- Vast majority of applications have the **bottleneck in the database not the code!**

Watch Your Errors

- Excessive, non-critical errors, such as E_NOTICE or E_STRICT can only be detected via error-logging.
- PHP code that generates any **errors** is going to **impact performance!**

Not Easily Detectable by Profilers

Micro Optimization

- Takes a long time
- Won't solve your performance issues
- Almost guaranteed to break something
- $\text{Cost} > \text{Reward}$

Speed vs Scale

- If you are planning for growth, scale is far more important than speed!
- Focus on scalability rather than speed, you can always increase scalable app, by simply adding more hardware.

Don't Re-invent the Wheel



Most attempts to make “faster” versions of native PHP functions using PHP code are silly exercises in futility.

Write Only Code

- Removing comments won't make code faster
- Neither will removal of whitespace
- Remember, you may need to debug that mess at some point ;-)
- Shorter code != Faster Code

Thank You!

Any Questions?

Slides @ www.ilia.ws