



Drupal Performance In A Nutshell

Shadi Alian
servit Open Source Solutions

About Me

- Shadi Alian
- 2 years JEE developer
- 1 year Drupal developer
- Free Lancing with Servit
- Teaching at The University Of Jordan



Hardware

- CPU
 - More cores the better.
- RAM
 - Lots of ram.
- Separate server for DB.
 - Drupal does lots of DB action.

LAMP Stack - Linux

- Linux
 - Use a stable distro (Debian based, Ubuntu...etc).
 - DO NOT USE A DESKTOP VERSION.
 - Use packaged applications, easier to update.

LAMP Stack - Apache

- Loads unnecessary modules
 - `mod_proxy`, `mod_php` ← not a typo.
- `MaxClients`: total available memory divided by average process size.
- Multiprocessing module
 - Worker: multiple child processes with many threads (use with caution).
 - `ThreadsPerChild`
 - Prefork: multiple child processes with one thread each.

LAMP Stack - PHP

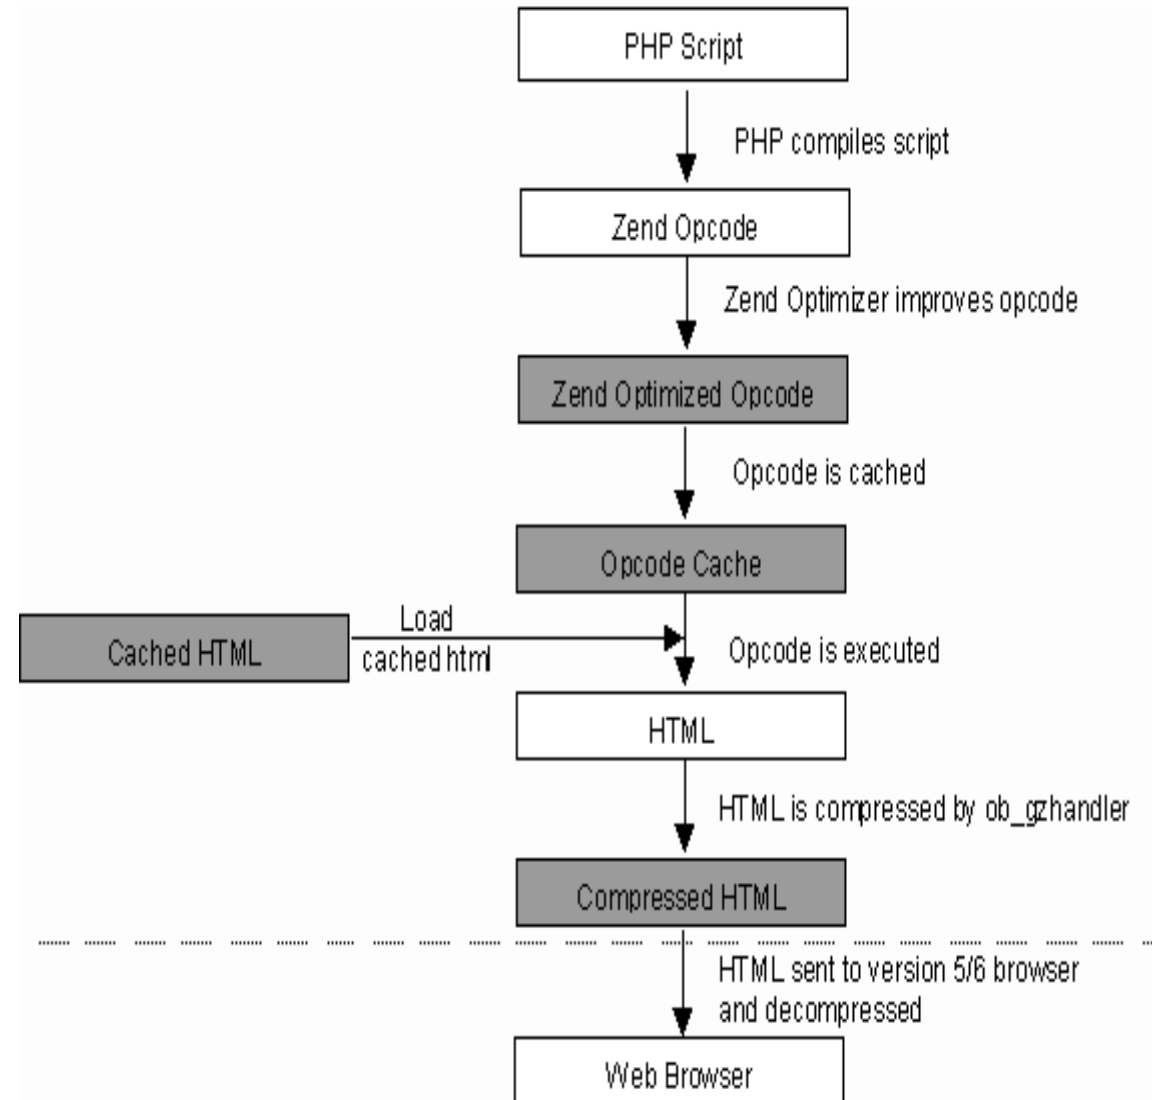
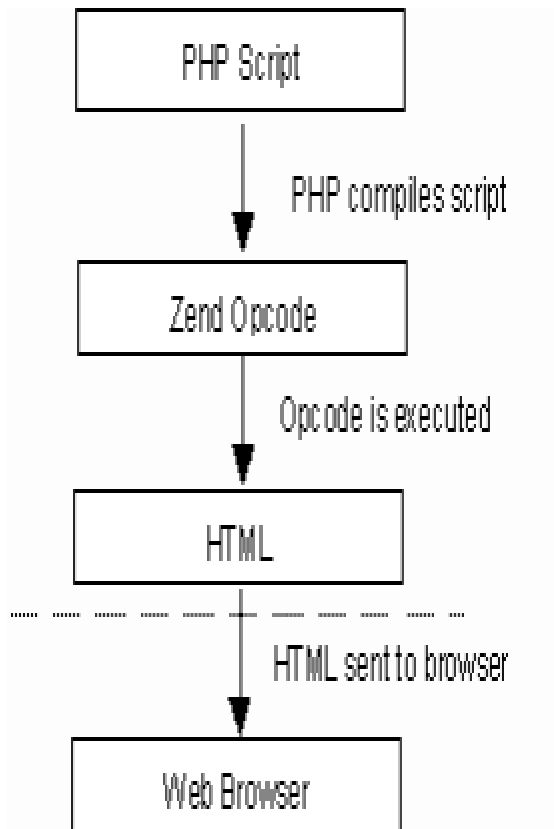
- `mod_php` default with apache.
- Install as FastCGI
 - Separate process
 - Less memory
 - Faster
 - Better to serve static pages
 - Hard to configure.

LAMP Stack - PHP

- Op-Code Caching
 - Less memory usage
 - Less CPU utilization
 - APC – open source
 - Zend optimizer
- Consider using HipHop
 - Built by Facebook (used in facebook)
 - Open source
 - Compiles php to optimized C++ code

LAMP Stack - PHP

- Op-code Caching



LAMP Stack - MySql

- InnoDB
 - Transactional
 - Better concurrency
 - innodb_buffer_pool_size 80% of memory
 - innodb_log_file_size 256M+
- query_cache_size
 - 32M – 512M

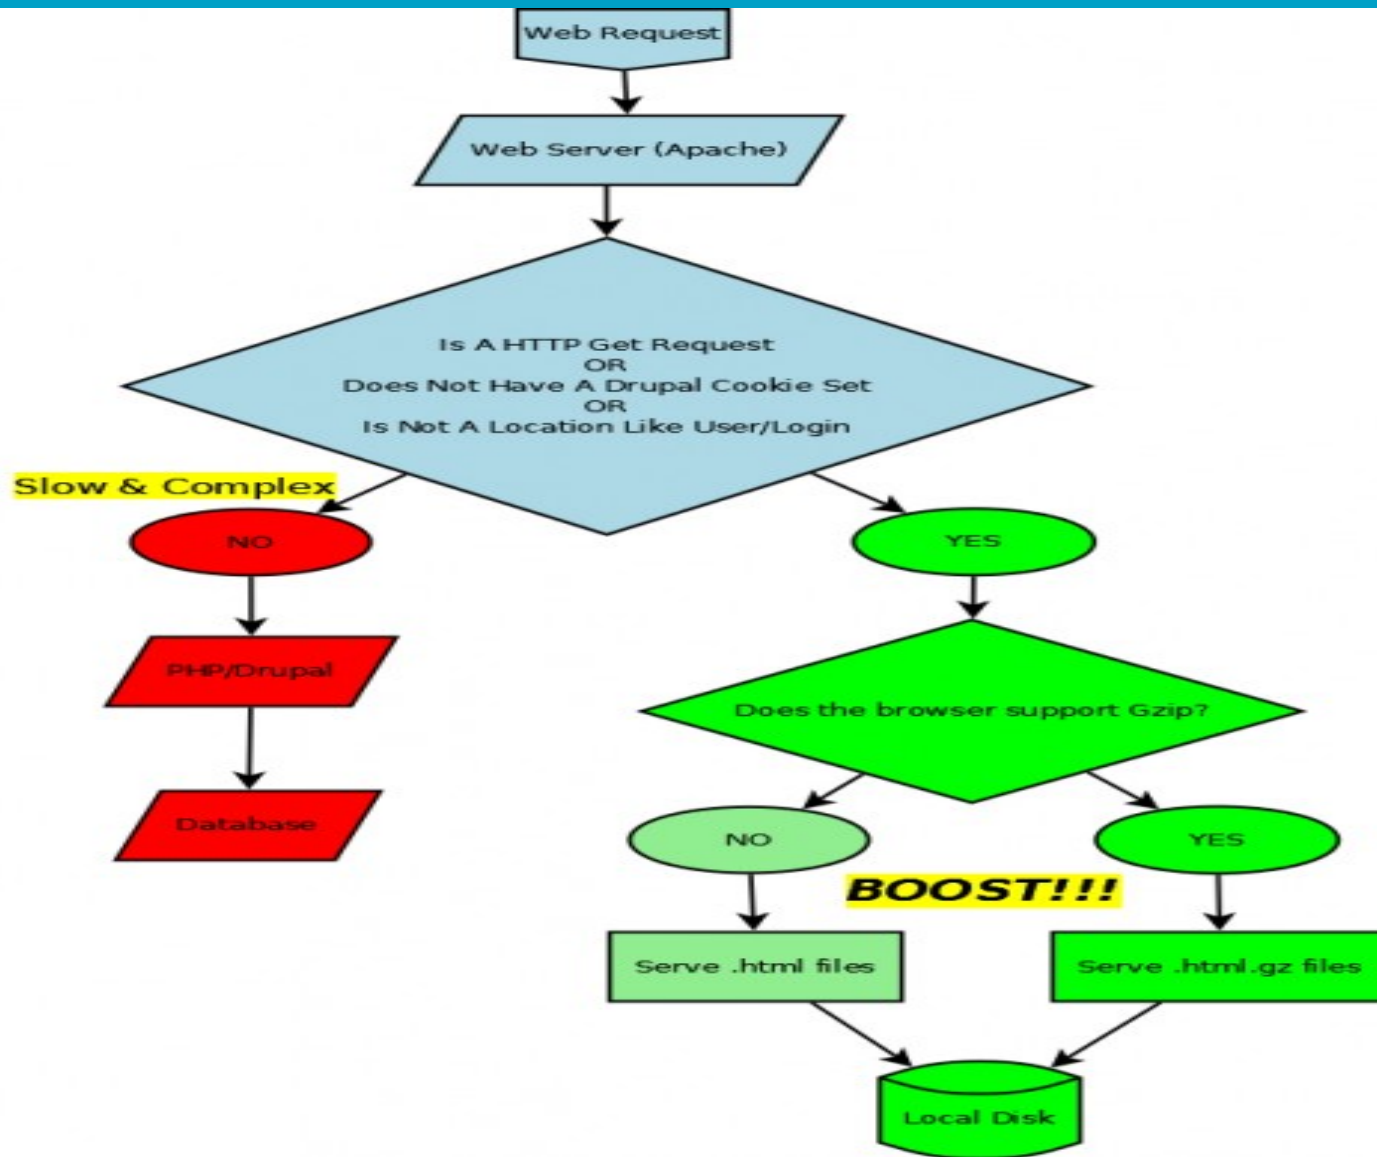
Drupal

- Memory intensive
 - D7 > D6 > D5
- Drupal tools
 - Devel
- Cache as much data as possible
 - Especially for anonymous users

Drupal contrib caching modules

- Static caching
 - Boost
 - Creates HTML for pages and stores it in files
 - Varnish has mostly replaced this module
 - Good for static pages and anonymous users data
 - Varnish
 - Reverse proxy to serve static files

Boost Preview



Drupal contrib caching modules

- Dynamic caching
 - EntityCache D7
 - Caches entities until expired or updated/deleted
 - APC
 - Caches objects in memory
 - Good only for single server
 - Memcache
 - Object caching for multiple servers
 - Slower on single servers than APC

Search Mechanism

- Drupal search works, but not as fast as outside solutions
- Searching the DB, Drupal's problem
- Solution: Apache Solr
 - Communication layer between the website and lucene search index.
 - Drupal module exist
 - Can be deployed on your system or used as an external service (Example provider: Aquia)

Other Options

- Pressflow
 - Drupal Distribution
 - Fully customized for performance and scalability
 - Varnish
 - APC
 - Can be built using HipHop - interesting
 - Many of improvements are in D7 core
 - 6 uses only MySQL
 - 7 is talking about using MongoDB

Thank you!