Crash-Only Software

George Candea and Armando Fox
Stanford University
{candea,fox}@cs.stanford.edu

Crash-only programs crash safely and recover quickly. There is only one way to stop such software—by crashing it—and only one way to bring it up—by initiating recovery. Crash-only systems are built from crash-only components, and the use of transparent component-level retries hides intra-system component crashes from end users.

What?

- Instead of:
  - Start-up
  - Orderly shutdown
  - Recovery (a special kind of startup)
- Just have:
  - Recovery

- Shutdown is either:
  - cut power
  - simple, external code (kill -9)

Why?

- Simpler, more robust
- Often faster

Within Components

1. All persistent(*) state in specialist state managers

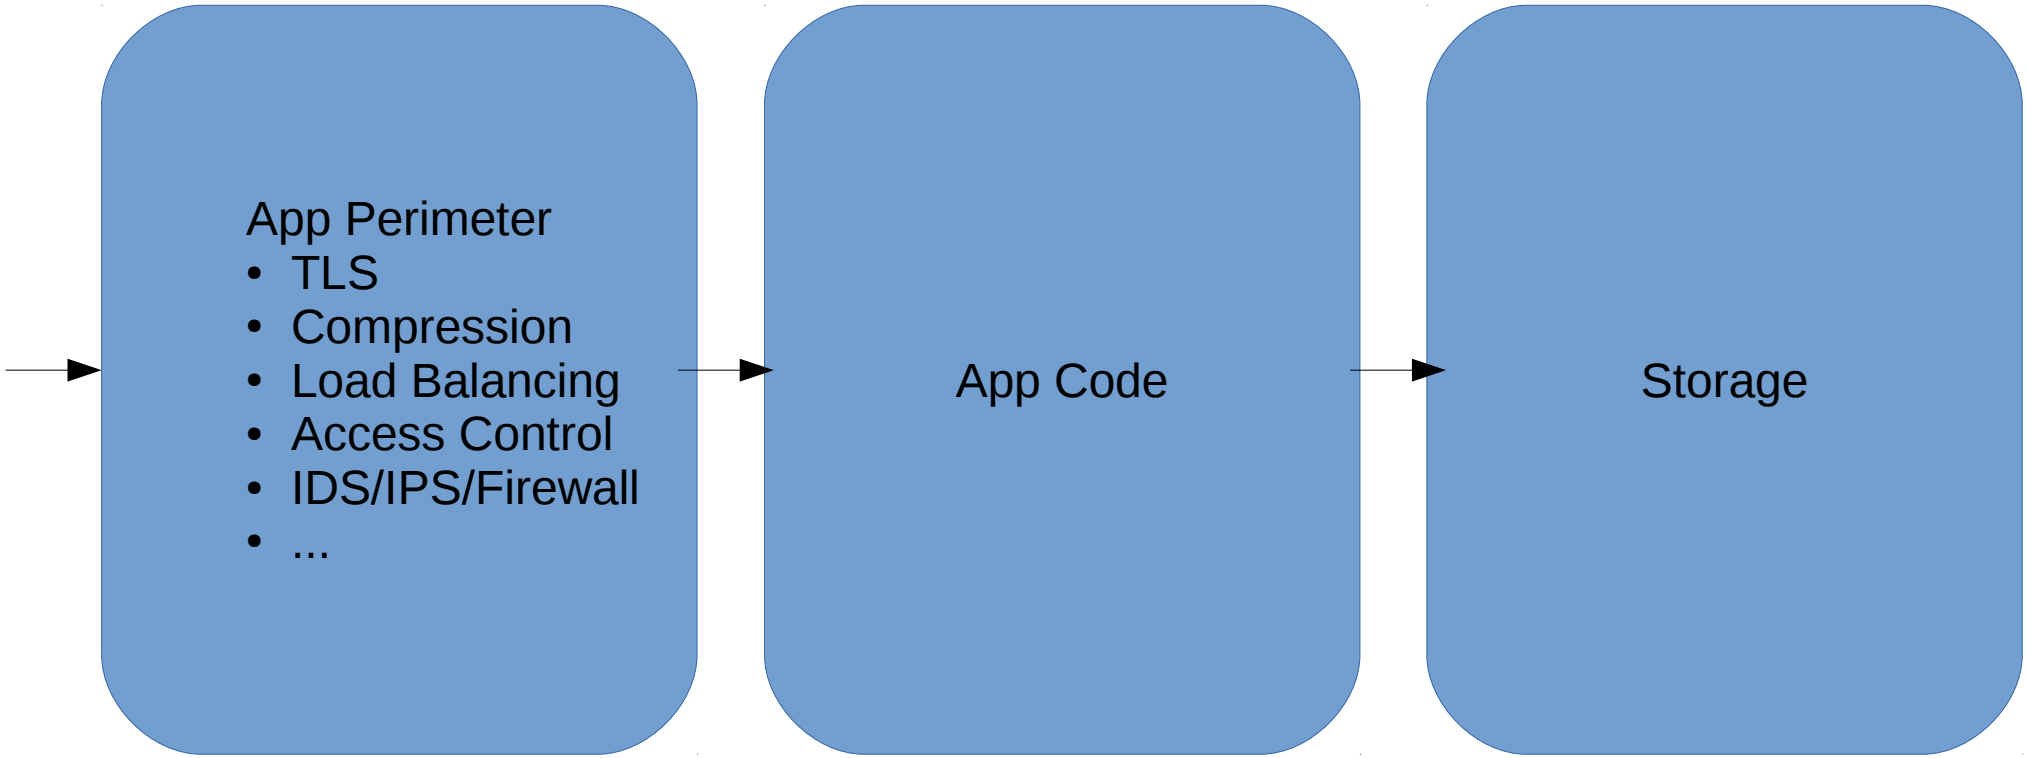

*: Typically: "between requests/operations"

2. State managers are crash-only / crash-safe

- Oracle vs. PostgreSQL

3. Use appropriate storage interface abstractions

- Minimise complexity (ACID / objects / session state / read-only / caches)
- Tie cost to requirements

- Don't use a filesystem for storing customers, opportunities, and orders
- Don't use an ACID database for caching

Between Components

4. Component boundaries contain failures

- VMs, Java processes/tasks, ...

5. All requests have timeouts

- Timeout triggers reporting to a recovery agent which crash-restarts the failed component
- Optionally: Client waits a while and retries a limited number of times
- Optionally: Failing component returns 503 and set "Retry-After: 1.3" header ; times spread out to soften restart
- Contains failures

6. All resources are leased

- Forces/allows reasoning about what happens when things go away

7. Requests are self-describing continuations

- (a data structure which describes execution state in a way that is programmatically accessible, rather than hidden in the runtime environment)
- Specified TTL
- Idempotent where possible (and labeled)
- Rollbacks only required where non-idempotent

Examples

- Chat daemon at UTS
- RealMail
- LPC
- Puppet vs. Debian
- Transactional database stores
- UNIXWare filesystem

Related Ideas

- Coerce failures to component failures ->

  all recovery is "restart failed component"
- Nightly reboot
- Beware Arianne 5 situation: recovery code is not optional!
- Simplifies rejuvenation (rotating VMs to deal with migration ; apply upgrades ; test for failures)
- Crash-Only “levels” ala RAID