
FUZZYLOGIC
www.fuzzy-logic.org
law@fuzzy-logic.org
206-432-6466

law@fuzzy-logic.org
Seattle, WA
206-432-6466

Fuzzy Logic Corp

Black Friday Readiness Assessment

law@fuzzy-logic.org
Seattle, WA
206-432-6466

How to Use This Self-Assessment

Here at Fuzzy Logic, we love eCommerce infrastructure. We've seen good patterns, anti-patterns, and novel implementations to thorny problems many times in our day. To help you prepare for what may be your business's biggest revenue event of the year, we've compiled this self-assessment to help you gauge your organization's readiness. These are not strict guidelines! They are merely talking-points meant to get you thinking about the common infrastructure and services many (but certainly not all) eCommerce companies rely on.

Each item has three levels, from 1 to 3, to assist you in gauging each area's maturity. If you find yourself with a good mix of 2's and 3's, rest assured you are ready for the Black Friday (and Cyber Monday) surge. If you're mostly 3's, congratulations are in order! However, if the self-assessment leaves you with some concerns or some difficult conversations to be had, feel free to give us a call. We're happy to discuss your readiness in detail, and how we may be able to help.

Finally, if you've got strong opinions one way or the other about this assessment, please feel free to reach out! We love learning what we got right and what we may have missed.

Yours in scale and stability,
--Lee Whalen
Founder and Principal Engineer
fuzzy-logic.org

FUZZY LOGIC

law@fuzzy-logic.org
Seattle, WA
206-432-6466

The Readiness Assessment

How much revenue do you lose per minute/hour of downtime?

L1: We don't know.
L2: We estimated it at $X a few quarters ago, and have run with it ever since.
L3: It's currently at $X, we review the estimate quarterly and update it as needed.

How easy/hard is it for you to add servers/capacity?

L1: Primarily 'by hand'; following a checklist or operator memory.
L2: Mostly automated; a VM is spun up and config-management is applied to bring it to spec.
L3: Fully automated; monitoring detects capacity is needed and deploys additional capacity without human intervention.

How much of your server configuration process is automated?

L1: None; each instance needs the bulk of its configuration applied by a human.
L2: Some; the heavy lifting is done via automation, but some hand-finishing is needed before adding to Production.
L3: All/most; human intervention is rare, and considered a 'bug' when it is needed.

What metrics are you capturing, graphing, and/or alerting on?

L1: Individual-server-based (CPU, memory, disk, etc), basic site availability (up or down?)
L2: The above, plus basic frontend/backend metrics (cache performance, search speed, web/app server performance, etc).
L3: The above, plus OKR/KPI's relevant to the business (# of sales/hr, # of searches/hr, etc).

How much of your code deploy process is automated?

L1: Very little; we ssh around to servers, 'git pull' to deploy, and reconcile configuration files manually after the fact.
L2: Some; code (either a framework or a for-loop) gets the code on the servers, with config-management handling the config files, but testing is an afterthought.
L3: All; we have a comprehensive CI pipeline that pushes to a 'staging' environment for automated testing. On acceptance, the code is seamlessly moved to Production. If there are issues at any point, roll-back is similarly easy/hands-off.

At what layers of the stack are you caching?

L1: What's a cache?
L2: There's some combination of frontend (CDN-based) and back-end (query or static-asset) caching, but little introspection on how efficiently it's performing.
L3: There's front- and back-end caching as dictated, with metrics under regular review to track and

FUZZY LOGIC

law@fuzzy-logic.org
Seattle, WA
206-432-6466

continuously improve cache efficiency.

How do you handle async jobs?

L1: They run on the same instances as our app servers, tightly coupled to the main application codebase.
L2: They run on dedicated 'cron' instances, at pre-set times (2x/day, 10x/day, every 2 minutes, etc).
L3: The jobs are factored into tasks suitable for distribution to 'queue' hosts so that capacity can be transparently added/removed as the business needs dictate.

Do you have redundant payment processing capability?

L1: No.
L2: Yes, we maintain a fully redundant PCI-DSS environment interfacing directly with banks/card providers.
L3: Yes, we outsource to specialized payment processing services (Stripe, GLBPay, etc) and we track the rates of success/failure for payments to each provider to calculate the reconciliation cost.

What is your alert escalation process?

L1: Customers alert us via email or social media.
L2: The ICs in charge of the services get alerted.
L3: We have a tiered on-call tree, escalating as high up the leadership chain as necessary to maintain a timely SLA.

How many single points of failure does your infrastructure have?

L1: We don't know.
L2: Greater than 2, but less than 4.
L3: 0-1 (and that '1' is really for 'we're in a single datacenter for the time being').

How complete is your runbook for common failures?

L1: What's a runbook?
L2: We have a few topics in there, but it's been some time since we've reviewed.
L3: Topics are added automatically as part of an issue's triage process, and disseminated throughout the team as part of a regular 'watch one, do one, teach one' good practice.

How many engineering-hours are spent manually reconciling failed transactions?

L1: We don't know.
L2: We have a dedicated department for that.
L3: Less than 5 hours a quarter.

law@fuzzy-logic.org
Seattle, WA
206-432-6466

Wrapup

Thanks so much for your time and attention. I would love to hear any questions/problems/cries of anguish you may have regarding the self-assessment - what works, what doesn't, what may be missing, etc. In that vein, please feel free to contact me any time at my email above. If you thoroughly enjoyed this, please don't hesitate to share it far and wide.