

DAVE JOSEPHSEN

iVoyeur: the dangers of marketing



Dave Josephsen is the author of *Building a Monitoring Infrastructure with Nagios* (Prentice Hall PTR, 2007) and senior systems engineer at DBG, Inc., where he maintains a gaggle of geographically dispersed server farms. He won LISA '04's Best Paper award for his co-authored work on spam mitigation, and he donates his spare time to the SourceMage GNU Linux Project.

dave-usenix@skeptech.org

I, like so many of my peers, have occasionally observed that if not for the presence of users, my systems and networks would perform perfectly. I suspect, however, that if we were ever to achieve our collective long-fancied dream of . . . “removing” the users, that running systems would be rather less fun, not to mention less profitable (but still worth it, I dare say). Though they are often worthy of my disdain, users are, I must admit, beyond necessary in my current line of work, where we live or die by the number of folks voluntarily using the Web sites we host; they are in fact . . . desirable.

That I should actually encourage people to use the services I'm providing smacks of poetic justice. At least it's not something my career to this point has prepared me for. Having been brought up mostly in large corporate environments, I've been of the mind-set that work gets done in spite of the users. Not even users, in fact, but lusers [1], pebcaks [2], and suits [3]. Users went from annoying to desirable in my estimation gradually. Truthfully, I still haven't entirely made the transition. It's been like slowly coming to the realization that carbs might not be as bad as everyone says. Intellectually I get it, but blueberry muffins still inspire pangs of guilt.

That users could be something to fear, however, was an entirely new and unconsidered possibility. Worse, I had no time to prepare mentally. I went into a meeting one day and came out a changed man upon being given the news that not only would a site we were hosting be given the front-page ad on msn.com, but that there would be several prime-time TV commercials during, for example, “ER,” “Lost,” and something called “The NCAA Playoffs.”

What kind of traffic increase can you expect to get from a TV ad? I did a quick mental inventory and discovered that I didn't know anybody who could easily answer that question. I've heard more than my share of slashdotting horror stories, but these anecdotes don't help, because nobody who tells them was capable of serving the traffic, much less measuring it. Alas, Google wasn't much help either. So in the hope that some terrified sysadmin finds it helpful, this month I offer to you my first experience with a real marketing behemoth.

I should back up a bit and give you some context. I left my quite cushy corporate job three years ago

to work at a tiny little tech company. Well, that isn't exactly what happened. I left under some duress, and wasn't getting along with my cushy corporate overlords. I had a distressing tendency to get things done, you see. Things that everyone wanted but nobody would dare request. So maddening was my penchant for implementing systems that the company needed yet management hadn't asked for that I was flung from manager to manager like a radioactive potato. I had 16 or so managers while I was there, and yet despite the re-orging, the time tracking, and even the Microsoft project, up would pop a monitoring system, a code promotion system, a documentation system. You get the idea. They were all free and they worked well enough that no one could stomach having them dismantled, but they didn't endear me to management. When the managers finally prevented me from getting anything done for nine months or so, I started sending out the resumes. When I gave notice two weeks later they were quite surprised, having been happy that I was finally settling into the corporate culture.

I was the first dedicated sysadmin-type in my current company. We have not many employees—fewer, in fact, than the number of Web pages we host. We have a great niche hosting “loyalty management” (points/miles etc.) programs for much larger companies. Because of the size of our customers compared to ourselves, the job works out to be what a lot of sysadmins might consider a dream job. Suffice it to say, no pigeonholing goes on here; servers, firewalls, routers, load balancers—the sysadmins (both of us) are armpits deep in all of it, and it's interesting work. The problems are large and varied, at least for me. On any given day I might, for instance, be designing BGP-based failover for our geographically dispersed co-located data centers, extending our LDAP schema, configuring LUNs on a SAN, creating VLANs, or writing bash scripts to transfer and load Oracle data.

The best part about this job is having the complete freedom to choose the right tools to do it. As a result we have quite an eclectic mix of tools, including OpenBSD routers on rather nifty Axiomtek-embedded hardware [4], Fujitsu SANs, various distributions of Linux including (shameless plug) SourceMage [5], running on everything from gray boxes to Sun x4550s, and, of course, the usual suspects that you'd expect from Cisco. Most of the best stuff we have is home-grown. I am especially proud of our load-balancer tier; it does things no commercial balancer can, and as much as I'd like to co-opt this article into fully describing it, I'll restrain myself and leave it to my cohort to publish, it being his brainchild. How's that LISA paper coming, Jer?

The problem with freedom is that pesky responsibility, so with the freedom of choosing and building your own tools comes the responsibility of building a solution bullet-proof enough to someday survive national TV and radio ads. By now I hope you can begin to relate to the sinking feeling I experienced in my stomach when given the details of our customer's marketing plans—plans that included not only a litany of TV and Web advertising, but 25% market saturation on a slew of national radio stations. Sure, we'd load-tested the environment, but at the time that thought provided little comfort. It's this feeling, or something like it, that I imagine made Microsoft so popular in the '90s.

I'll cut to the chase here and tell you that we've survived the onslaught (so far), and that the marketing was able to increase our traffic by more than an order of magnitude. Figure 1 is a graph of the traffic during the initial marketing ramp-up. The y-axis has been altered to protect our customer, but otherwise the data is unaltered.




FIGURE 1

The point marked A shows our normal pre-marketing traffic. Point B marks the beginning of the Web marketing campaign, which included the aforementioned msn.com front-page ad, and point C marks the beginning of the radio and TV ads, the first of which appeared during a popular prime-time ABC show.

We used the week or so of lead-time we were graciously given to scale up our app tier. This was done easily enough, because the app tier servers are Xen virtual machines. We quickly doubled their number, and distributed the app across them. We also built up some additional balancers in case we needed to scale up the balancer cloud, but as of this writing we haven't had to use them. Quite to our relief, the traffic increase was pretty much ideal from our perspective, *lots* of extra traffic but nothing the systems couldn't handle.

What about security? Does being on TV make you a bigger target? To that question I can give an emphatic yes. And while I can't give a whole lot of detail here, I can say that application-layer attacks (injections, URL traversal, cross-site whatevertheheck, etc.) have increased nearly 600%. The lower-level stuff—scans, brute-force attacks, etc.—have increased in kind. The cycles we spend mitigating attacks are not trivial. Were I advising someone building an infrastructure from scratch, my advice would be to make attack mitigation, logging, etc., a design consideration. But that sort of goes without saying nowadays.

For now at least, it seems I've reached a happy medium with my users, though I suspect I remain more frightened of them than they of me. That particular role reversal is troubling, to be sure, but worth it, I dare say.

Take it easy.

REFERENCES

- [1] <http://catb.org/jargon/html/L/luser.html>.
- [2] <http://catb.org/jargon/html/P/PEBKAC.html>.
- [3] <http://catb.org/jargon/html/S/suit.html>.
- [4] <http://www.axiomtek.com.tw/Products/ViewProduct.asp?view=676>.
- [5] <http://www.sourcemap.org/>.