

/var/log/manager

When Technology Isn't the Cause of a Technical Problem

ANDY SEELY



Andy Seely is the manager of an IT engineering division, customer-site chief engineer, and a computer science instructor for the University of Maryland University College. His wife Heather is his init process and his sons Marek and Ivo are always on the run queue. andy@yankeetown.com

If the only tool you have is a hammer, every problem looks like a nail. In the technology space, we tend to approach all problems as technology problems. That's how we're wired. We're systems people. If something's not performing correctly, maybe we can adjust the system settings or the resource provisioning. Maybe we can buy a new software tool to compensate. It's a different kind of management challenge to see a technical problem's organizational roots and to make an adjustment, far removed from the actual technology, that can relieve tension in the organization and result in better system performance.

It's Just a Slow "Picard vs. Saruman" Sort of Day

Everything's fine. The enterprise is running within parameters. The technical team seems happy, at least in that they're not dealing with things any more important than Pickard vs. Kirk (see Sidebar). The management team's biggest problem is worrying about which Web sites the employees are surfing to when they're on the clock. Customer calls to the service desk are normal noise: password resets and unreasonable demands for magical computers that don't exist. It's a good day.

We use an internal service to process requests for a business intelligence (BI) product. It's a pretty sizable data warehouse with a number-cruncher front end. There's a small team of operators and a couple of sysadmins, all of whom keep mostly to themselves. It runs; no one worries. Once the Picard vs. Kirk got some Saruman and Gandalf thrown into the fight (without Picard and Kirk exiting, which is in itself interesting), I wandered down the hall to ask how those keep-to-themselves BI folks were doing.

Everything's fine. Customer queries were being answered. They had no problems. Except that performance wasn't really what they liked. OK, performance tuning is something we do, so I asked them to describe the performance problem. Well, they say, the front end has been broken for months, and the sysadmin can't keep up with the operator requests that he's been answering directly from the database using the command line SQL interface.

What?

Finding Common Ground in a "Picard vs. Saruman" Sort of Situation

I'm not a fan of meetings for the sake of meetings, but if ever there was a need to get everyone around the same table, this was it. We called a meeting of the BI team, the engineering team, the storage team, the operations team, the database team, the monitoring team, the service desk, and the management team.

From the start, we didn't have consensus. Each area of our overall team felt that they either already knew the whole story or didn't have any responsibility for this system at all. I like to joke that I'm "classically trained" in the art of holding meetings, but this was a tough one to navigate. I'm the head of engineering and have a finger on the pulse of almost all we do, but this Business Intelligence system pre-dates me, and almost no engineering project or sup-

`/var/log/manager`: When Technology Isn't the Cause of a Technical Problem

port work had come up during my tenure. The meeting was as much about discovery for me as anything.

And what did I discover? Operations had no sense of ownership because the BI team had their own internal sysadmins. The BI sysadmins felt like they had been abandoned, because they had submitted dozens of service requests to operations over the years, and they couldn't understand why no one realized they had a systemic problem. The engineers acknowledged the presence of the BI system but only so much as they occasionally got asked for very specific help doing very specific tasks. The storage team responded to storage requests like they do to everyone, with an initial "we don't have any more storage capacity" followed by a grant of storage after they discovered a way to free more space. If we'd all had pistols, we'd have looked like a Spaghetti Western, with all of us pointing our guns at each other.

Leadership with a Lowercase "l"

I'm not the manager of the BI system or of the operations team, but I am a senior manager in the technical staff. If I see something broken, it's my responsibility to ensure it gets fixed. There are many schools of thought on leadership. I chose to employ my own "big-L/little-l method." This meeting cried out for some "little-l," or lowercase leadership: It didn't need some big boss to make big decisions, just someone to get his hands dirty and help clear the path so that everyone could have a say and get all the facts on the table.

I guided the discussion and turned it over and around until everyone at the table had the same basic understanding of the BI system architecture and dependencies. Then we drew it on a whiteboard and walked through it again, refining the diagram until it reflected both the system and our common understanding.

After we all agreed on architecture, we walked through data flow. Request comes in, gets received here, gets processed here, traverses this subsystem and that subsystem; we followed the flow from query to answer. We talked about system failures and how they're reported and recovered. We talked about resource provisioning and network link speeds. We asked the functional expert to talk about the BI system's internal limitations for complex queries and how the vendor's tuning recommendations were being applied.

Leadership with an Uppercase "L"

We discovered some non-obvious but fundamental flaws in the system, but not the system one would think. Our technical flaws were coming from the organization itself.

1. The BI team gave the appearance of running their own show. The operations team didn't track metrics or report BI outages on their balance sheet, which meant that operations management never put pressure on the BI system to be tuned or improved from a systems perspective.
2. By having its own sysadmins, the BI team built an unintentional wall between themselves and the rest of the sysadmin team. The operations sysadmins never added up all the little requests for support to make a bigger-picture approach because they figured the BI sysadmins knew what they were doing. By reporting issues through business rather than technical management chains, the BI sysadmins' complaints up their management chain fell on deaf ears.
3. By not having storage engineers involved with a holistic perspective, requests were fulfilled as requested rather than as needed, and they weren't requested in such a way as to put database indexes on the fast storage.

Fixing this required "big-L," or uppercase Leadership: The boss needs to make changes in how we do business.

True Story: The DNS Subdomain Generation and Genre Problem

We were troubleshooting a DNS problem with a delegated subdomain. When we started looking into the architecture of the subordinate organization, we found that they had four redundant DNS hosts with host names "picard," "kirk," "gandalf," and "saruman." I was leading the technical team researching the problem. We found the root cause was bad glue records, but in my final analysis I pointed out that there were at least two major system incompatibilities in the subdomain. First, there's a generational gap; Picard and Kirk are not going to cheerfully serve up the same answer as peers. Kirk will overpower Picard whenever he can, serving DNS answers that are the best for Kirk's own position. Second, there's a genre gap; you can't have Kirk and Gandalf working in the same DNS namespace. You'll get DNS query responses in Elvish one time and in Klingon the next, obviously resulting in protocol errors. Our recommendation was to rebuild the whole DNS environment and rename with more of a modern meme. The DNS servers should be: "neo," "morpheus," "trinity," and "tank." This way, they're all on the same team, serving the same mission. Performance will be improved through the virtualization of three DNS hosts, but it's a good idea to keep one DNS server physical to remove the common dependency on the virtual environment.

Starship Captains to the Bridge, Wizards to the Tower: Small Personnel Adjustments Can Make a Big Difference

We drastically improved the BI system with some small organizational changes. The BI sysadmins were reassigned to the application support group in operations. The enterprise monitoring team was given the green light to dig deeper and monitor more aspects of the system and to treat problems with more vigor than just sending an email to the BI team. The storage engineers were given greater purview over the “why” as well as the “what” when it came to decisions on storage provisioning for the BI system.

These organizational changes allowed real system improvements to flow:

1. Network links were upgraded and made standard between all BI systems, removing inter-system bottlenecks.
2. Database indexes were moved off SATA and onto solid-state drives, removing the BI query bottleneck.

3. Benchmarks were established for BI queries, creating a measuring stick of how to interpret BI system performance.
4. New monitoring hooks were established and alert playbooks created, improving overall awareness and problem response times.

These system improvements allowed the real benefit to happen: The BI query backlog was eliminated, and the BI functional operators were able to do their own jobs effectively. Getting there wasn't obvious, and it took a combination of the little-l leadership of guiding people to talk to other people and the big-L leadership of making immediate organizational changes in multiple areas to get work flowing and the system back to its core function of making money for the company. A reorganization of the team wasn't the most obvious approach, but ultimately it was the correct one. I'm the manager. That's my job.

USENIX Member Benefits

Members of the USENIX Association receive the following benefits:

Free subscription to *login*, the Association's magazine, published six times a year, featuring technical articles, system administration articles, tips and techniques, practical columns on such topics as security, Perl, networks, and operating systems, book reviews, and reports of sessions at USENIX conferences.

Access to *login*: online from December 1997 to this month: www.usenix.org/publications/login/

Access to videos from USENIX events in the first six months after the event: www.usenix.org/publications/multimedia/

Discounts on registration fees for all USENIX conferences.

Special discounts on a variety of products, books, software, and periodicals: www.usenix.org/membership/specialdisc.html.

The right to vote on matters affecting the Association, its bylaws, and election of its directors and officers.

For more information regarding membership or benefits, please see www.usenix.org/membership/ or contact office@usenix.org. Phone: 510-528-8649

USENIX Board of Directors

Communicate directly with the USENIX Board of Directors by writing to board@usenix.org.

PRESIDENT

Brian Noble, *University of Michigan*
noble@usenix.org

VICE PRESIDENT

John Arrasjid, *VMware*
johna@usenix.org

SECRETARY

Carolyn Rowland, *National Institute of Standards and Technology (NIST)*
carolyn@usenix.org

TREASURER

Kurt Opsahl, *Electronic Frontier Foundation*
kurt@usenix.org

DIRECTORS

David Blank-Edelman, *Northeastern University*
dnb@usenix.org

Cat Allman, *Google*
cat@usenix.org

Daniel V. Klein, *Google*
dan.klein@usenix.org

Hakim Weatherspoon, *Cornell University*
hakim@usenix.org

EXECUTIVE DIRECTOR

Casey Henderson
casey@usenix.org

SAVE THE DATE!



11th USENIX Symposium
on Operating Systems Design
and Implementation

October 6–8, 2014
Broomfield, CO

Join us in Broomfield, CO, October 6–8, 2014, for the **11th USENIX Symposium on Operating Systems Design and Implementation (OSDI '14)**. The Symposium brings together professionals from academic and industrial backgrounds in what has become a premier forum for discussing the design, implementation, and implications of systems software.

Don't miss the co-located workshops on Sunday, October 5

Diversity '14: 2014 Workshop on Supporting Diversity in Systems Research

HotDep '14: 10th Workshop on Hot Topics in Dependable Systems

HotPower '14: 6th Workshop on Power-Aware Computing and Systems

INFLOW '14: 2nd Workshop on Interactions of NVM/Flash with Operating Systems and Workloads

TRIOS '14: 2014 Conference on Timely Results in Operating Systems

All events will take place at the Omni Interlocken Resort



www.usenix.org/osdi14