ALVA COUCH

should the root prompt require a road test?

MUSINGS FROM A SYSTEM ADMINISTRATION TEACHER ON THE IMPORTANCE OF CAREFULLY OBSERVING BEHAVIOR, AND THE TRUE COST OF CERTIFYING A SYSTEM ADMINISTRATOR

Alva Couch is an Associate Professor of Computer Science at Tufts University, where he and his students study the theory and practice of network and system administration. He served as Program Chair of LISA '02 and was a recipient of the 2003 SAGE Professional Service Award for contributions to the theory of system administration. He currently serves as Secretary of the USENIX Board of Directors.

couch@cs.tufts.edu

HOW DOES ONE HIRE A COMPETENT system administrator? As one who has taught system administration at the college level a small number of times in the past few years, I cannot help but wonder what people are thinking when they hire a "certified" system administrator over an "uncertified" one. With a small number of exceptions, most certification programs require only passing a written test to become certified. In teaching system administration, I have given many written tests to my students, and I can testify from personal experience that it is possible to pass a written test on system administration and not have the slightest clue about how to function as a system administrator.

Because I possessed a "test form reader" that could read standard test forms, and because I possess in good measure Larry Wall's "attributes of a good programmer" including laziness, impatience, and hubris, I was strongly motivated to make multiple-choice questions "work" to evaluate students so that the card reader could do the work of grading for me. Thus I went to great lengths to make multiple-choice questions "difficult enough" to demonstrate system administration knowledge instead of luck. My early efforts included a test with 10 possible answers per question, and a test with 100 true/false questions, where with each incorrect answer a point gets subtracted from one's score, so that the "random chance" score is 0. I know of at least one student who passed each of these tests with flying colors whose incompetence as root became clear when one *observed the person's behavior*. So those tests were out of the running as an effective measure of system administration competence. The form reader went into the trash.

Currently, I use fill-in-the-blanks questions rather than multiple-choice questions, to give examinees much more creative ways to make mistakes. This takes a little more time to grade, but at least I hope that the student who passed the multiple-choice tests above without knowing anything would not have a chance of answering fill-in-the-blank questions correctly. Still, I would not give an examinee with a perfect score on these tests access to root on *my* systems.

Instead, I might watch a potential administrator perform some tasks and rate how well the person navigated and managed the complexity of the tasks, how the person's own knowledge gaps were filled, and how many false starts and poor practices occurred between start and finish of each task. How can we put this "watching" into a reusable container that others can use as well? The answer may lie in studying how other fields certify their practitioners.

Licensing Drivers

My experiences as tester seem to indicate that no matter how carefully or well a written test is constructed, one cannot expect that a person who passes a written test on system administration can "take the driver's seat" without problems, any more than one can safely put a teenager in control of a car solely based upon his or her scores on the written exam. This leads me to consider parallels between driving tests and system administration certifications that might guide us to a better understanding of certification and what different kinds of certification might mean.

There are three components to obtaining a typical driver's license as a teenager:

- Experience: Has one read the rules and sat behind the wheel for a while?
- Knowledge: Is one aware of the laws and regulations?
- Function: Can one control a car in a realistic situation?

A typical teenager gains experience and knowledge from a driver training course and takes two tests: a written test of knowledge and a road test of driving under realistic conditions. A universal aspect of driver (and boat or pilot) licensing is that *passing the written test entitles one to take the road test*, and nothing more.

Following the parallels between driver licensing and system administration, *current written test certifications seem to entitle the certified person to a "road test,"* and nothing more. Employers must perform the road test themselves, the hard way, by putting "certified" individuals at the root prompt and observing in a live situation whether problems develop. This is like putting the teenager who has passed the written exam at the wheel, and then taking the license away only after the first pedestrian is run over.

What, then, comprises a good system administration "road test"? We can learn again from the components of typical driving road tests; one has to be able to function in a variety of common situations and make good decisions in those situations. But there is another component to the driving test that is often ignored: There is an *observer* who evaluates factors other than just being able to perform the tasks requested. Confidence, attitude, and situational awareness are being observed as well. The driver's test is not scored solely on whether one can navigate from point A to point B without accident, but rather on how one approaches the problem and handles unexpected events.

Situational Awareness

Here, human factors research can help us understand the problem of "driving" either automobiles or configuration changes. The "situational awareness" component of driving, flying, and even managing systems has been extensively studied, mostly to determine whether controls and cockpit layout aid or hinder the driver or pilot in performing a task. In human terms, the expert driver exhibits constant eye movement, checks constantly beside and behind the car when driving, and maintains a "situational awareness" that helps the driver react quickly to contingencies. This awareness can be tested, but only through direct observation of someone driving, and not simply by observing successful navigation from point A to point B.

As an example, let us consider two system administrators engaged in a road test. Each system administrator is asked to configure security for an Internet service. The first system administrator (X) tries to edit an appropriate file, realizes that it is not writeable to his or her user, su's to root, makes the file world-writeable (chmod 777), edits the file, and then makes the file normally writeable (chmod 644), without checking the prior protections of the file. The second system administrator (Y) lists the file to check protections, looks up the appropriate format for the change in the documentation, su's to root, edits the file, and exits the root session. Which of these system administrators should pass the road test?

Passing the first system administrator (X) on this road test is like passing a new driver on a road test in which the driver smashes the bumper of a parked car when parking, but pays the owner in full for the damage and replaces the owner's bumper with one of his or her own choosing without referring to the model of car that was hit. Further, during the road test, the driver leaves the car on the highway with the motor running, steps out to shop (during which time anyone could have gotten behind the wheel and driven away), comes back as if nothing has happened, and considers this all to be good driving practice.

The main point of this simple example is that the road test should measure process rather than product. It is not sufficient that system administrator X successfully moved from point A to point B, given that the path taken from A to B was problematic and irresponsible. But, to determine this, the examiner must have a high level of situational awareness, so that it will be possible to determine when mistakes are being made. Finally, the examiner must be an expert and trained observer in order for the road test to be a reasonable measure of prowess. Speed is not as important as safety and wise choices.

There are a spectrum of tools available to certify system administrators, from purely written tests to "live tests" in which a problem must be diagnosed and repaired. But the preceding example shows a situation in which simply accomplishing an action is not enough; the administrator must also adhere to good practice throughout, in action, word, and deed. To the best of my knowledge, no certification program—including even the highly acclaimed RedHat program that does require a live skill test—requires a stringently observed road test of this kind, in which the person to be tested is watched carefully and evaluated on process by an expert human observer.

I hope that potential employers understand this and conduct road tests of their own before allowing a "certified" system administrator free rein over their systems.

The burning question that arises is whether we *can* or *should* provide some external neutral mechanism by which system administrator "road tests" can be performed, or whether such "road tests" are forever the responsibility of potential employers. In seeking a neutral, reusable "vehicle" for road tests, one can seek inspiration in the way that driver licensing is made reusable and widely applicable.

Measuring Trust

A second lesson we can learn from driver licensing is how one progresses from the basic certification to being certified at higher levels. The federal government in the United States mandates a "class system" in which there are three main classes of commercial drivers, defined by the weight or passengers of the vehicle to be controlled:

- Class A: Any combination of vehicles with a gross vehicle weight rating (GVWR) of 26,001 or more pounds, provided the GVWR of the vehicle(s) being towed is in excess of 10,000 pounds
- Class B: Any single vehicle with a GVWR of 26,001 or more pounds, or any such vehicle towing a vehicle not in excess of 10,000 pounds GVWR
- Class C: Any single vehicle, or combination of vehicles, that does not meet the definition of Class A or Class B, but is either designed to transport 16 or more passengers, including the driver, or is placarded for hazardous materials

As well, there are two noncommercial classes of license:

- Class D: Regular noncommercial vehicles
- Class M: Motorcycles

Finally, there are "endorsements" that one must obtain in order to drive under special conditions:

- T—Double/Triple Trailers (knowledge test only)
- P—Passenger (knowledge and skills tests)
- N—Tank Vehicle (knowledge test only)
- H—Hazardous Materials (knowledge test only)
- S—School Buses (knowledge and skills tests) [1]

The license itself is only part of the picture: The general license class is a measure of knowledge, skill, and trust, whereas the endorsements are indications of knowledge and skill.

This system is not based upon measuring capabilities of drivers, but, rather, upon the kind of vehicle to be controlled and the kinds of risks that must be mitigated in controlling that kind of vehicle. This practice inspires me to rethink how we might categorize certifications for system administrators. Almost exclusively, we tend to think of certifications as proving proficiency for specific platforms or products or even specific courses of study. A "certification to configure Cisco switches" seems—in light of this discussion—to be similar in character to a "license to drive BMWs," whereas a "certification" based upon taking a specific course in networking seems similar to a "license to drive in Boston." Both of these kinds of "certifications" look more like "endorsements," that is, something to be attached to an overarching "license" to indicate specific capabilities.

Simplicity

Another thing one can learn from the driver's road test is simplicity: Several basic skills are tested and nothing more. The reason for the simplicity of road tests is straightforward: Road tests are expensive to administer and thus must be made generic to be cost-effective. Is system administration any different? If it is the same, then what categories can we suggest for licensing and testing, and what can we learn from the structure of driver testing that can help us? Let us try a simple thought experiment and consider how system administration certifications might be categorized by risk class and what endorsements might apply. This gives rise to a very different design for certification than what is available now.

Let us consider one model of coarse risk classes that might be the basis for a new model of licensing and certification:

- Class D: Individual workstation, including multiple workstations that are part of some LAN or enterprise
- Class C: Server, including the potential that multiple servers are administered
- Class B: LAN, including basic automation of workstation and server configuration
- Class A: Enterprise, including enterprise-wide design and implementation

I do not defend this as definitive; it is just a "straw dog" proposal for how a "license-based" scheme of certification might be structured. The key ingredient of this system is that the class of a certification is based upon the risk entailed in managing that kind of system or network, and not upon a capabilities model of the administrator. The question is not whether a driver is *capable* of driving a larger vehicle, but whether she or he can be trusted to do so both safely and responsibly.

In turn, each system administrator "road test" would measure whether a person can be "trusted" with the responsibilities of that level of management. This is not a platform-specific basis for trust, even though someone might well be tested in the context of a specific platform. What is being tested is not the knowledge of the platform, but, rather:

- How the candidate responds to contingencies
- Whether the candidate adopts responsible practices in making changes
- Whether the candidate is aware of the effects of his or her actions
- Whether the candidate has appropriate knowledge-acquisition skills and understands personal limits
- Whether the candidate has effective interpersonal skills
- Whether the candidate is aware of and acts in accordance with the legal, ethical, and professional obligations of the system administrator

In other words, an administrator who has complete mastery of Linux, but who obviously does not follow the code of ethics when it is more convenient to ignore it, should fail the test. This is analogous to the driver's test in which the candidate has exceptional control of a vehicle but drives on sidewalks where possible.

The second component of a license-based model is a set of "endorsements" that might be obtained to operate networks in subclasses of the classes just listed:

- Directory services
- Electronic mail and spam
- File service, SAN, NAS, etc.

Note that the currently available "certifications" have become "endorsements," and many of these do *not* require a live skill test.

The risk model derived from motor vehicle operator licensing seems to imply that we are missing an overarching level of certification—involving some form of road test—before and not as a substitute for the specialty certifications that system administrators currently pursue. This taxonomy suggests to me that current certifications are valuable only in certifying skills of those system administrators who "already know how to drive" and—in some sense—are already worthy of some form of "driver's license."

Driver Training

A world full of licensed drivers who learn how to drive by causing accidents is a frightening thought, but it seems that this is exactly what we have in training system administrators. The lack of a road test that examines the most basic of skills, combined with certifications that measure only advanced knowledge, creates a knowledge gap that only experience fills. This hurts our image as professionals and makes the path to proficiency both haphazard and painful.

So what can be done to address this? I have no magic solutions, but I can suggest some strategies that might greatly improve training in the profession. Some of these are quite controversial, and I expect it would take years for us to agree on some of these points, but I will state them anyway:

- *Emphasize understanding of effect* of an action rather than just "what action to perform." Aim for situation awareness rather than mastery of recipes.
- *Replace "best practices" with "responsible practices"* as the most basic form of system administration education. In other words, emphasize responsibility rather than optimality in training beginners.
- *Emphasize professionalism and appropriate behavior* before skill and knowledge.
- *Emphasize experience as part of certification.* Just as truck drivers have to have a certain level of driving experience before they can achieve higher license classes, experience is underrated as a certification element.

To some extent, we are already doing this, but to some extent, we are also failing. Every employer has to devise a custom road test, or suffer the consequences. Every system administrator has to learn the hard way. "Road tests" (and the training that comes with them) are a distant pipe-dream that no one knows how to implement.

But there is one concrete and unavoidable conclusion from this essay that I do not believe is controversial: *Current certifications test knowledge and not professional practice*. If we develop a road test for system administration, it will not simply be about accomplishing tasks, but about behaving in a professional way in reacting to contingencies and requests. It will certify safety, responsible behavior, and situational awareness rather than demonstrating knowledge of how to configure Microsoft Windows. The latter is an endorsement that can best be earned when the candidate already has a firm foundation in professional practice.

This is not to say that current certifications do not have value. I am afraid, however, that the value that they represent is often misinterpreted or taken out of context. There is value in the knowledge of how to drive an 18-wheeler, but without a road test, there is no confidence that the person can be trusted to utilize that knowledge wisely. Change "18-wheeler" to "enterprise network" to obtain a clear picture of the crisis at hand.

Nothing springs into existence from nowhere without some form of evolution. We are "evolving" toward "road tests" at an alarming rate, in the same way that traffic accidents cause stoplights to spring up at busy intersections. It might be time for system administrators to rise to the occasion and fill the certification gap themselves, before someone else thinks of doing it for us.

REFERENCES

[1] Massachusetts Driver's License Web site.