
# Hard Drive Unreliability

**Description:** Leo and Steve discuss the distressing results and implications of two recent very large population studies (more than 100,000 drives each) of hard drive field failures. Google and Carnegie Mellon University (CMU) both conducted and submitted studies for the recent 5th USENIX conference on File and Storage Technologies.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-081.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-081-lq.mp3

---

INTRO: Netcasts you love, from people you trust. This is TWiT.

**Leo Laporte:** Bandwidth for Security Now! is provided by AOL Radio at AOL.com/podcasting.

This is Security Now! with Steve Gibson, Episode 81 for March 1, 2007: Hard Drive Reliability.

Security Now! is brought to you by Astaro, makers of the Astaro Security Gateway, on the web at www.astaro.com. And by Nerds On Site, helping people with technology all over the world. Visit Iwanttobeanerd.com.

This is the Security Now! I've been waiting for for a while. Leo Laporte here. Steve Gibson is in his laboratory watching hard drives spin.

**Steve Gibson:** That'd be SpinRite.

**Leo:** SpinRite, they're spinning right there.

**Steve:** Spinning right and spinning wrong.

**Leo:** Have you turned off all the noisemakers?

**Steve:** Ready to go.

**Leo:** Someday I want you to leave them on. That'd be fun.

**Steve:** You'd be able to hear Fred Flintstone say "Yabba dabba do" every so often in the background.

**Leo:** That's when you get a new registration for SpinRite and different sounds for different people's emails and...

**Steve:** Ah, yes.

**Leo:** Do you have any really rude sounds for anybody?

**Steve:** No, I've got some friends who have totally gone sound happy. I'll be talking to them on the phone, and it's sort of like people who have disco music for their ringtone on their phone. It's like, oh, what happened to a simple little chime?

**Leo:** You know what, after a long time playing with ringtones I finally went to the "nostalgia" ringtone which is basically a phone ringing. And it's not so bad. And nobody looks at me funny when my phone rings.

**Steve:** Exactly.

**Leo:** Last thing I had was the Call For Help theme. That really put it over the top. That was it.

**Steve:** When the phone starts doing a cha-cha all by itself, and then keeps on going. I mean, you know, meanwhile some embarrassed person is fumbling through their briefcase or through her purse or something because it's like, okay, well, we all know.

**Leo:** Well, if it's on the table and the vibrate's on, it might quite literally be doing the cha-cha right across the table. I've seen that happen, too. Hey, we want to welcome a new sponsor to the show. We're very glad to have them. You had coffee with them when you were up in Toronto last time. They're worldwide.

**Steve:** They have been Security Now! followers from the beginning, Leo.

**Leo:** I'm so glad. Their name is Nerds On Site, don't laugh, Nerds On Site. And their specialty is helping people with their technologies. Now, in the past you've thought, oh, I would like to get a consultant to come in, somebody who would help me troubleshoot my system or help me configure it or set it up. But you've wondered, how can I know whether they're any good or not? And that's really kind of the problem. That's why you're going to love Nerds On Site. All Nerds on Nerds On Site are independent contractors, but they're certified, and they have a University of Nerdology. And in fact they're looking for new Nerds. So that's one of the reasons they wanted to buy some ads, because they want to

tell you that, if you want to be a nerd, you can. Nerds On Site is in Canada, U.S., Mexico, England, Australia, South Africa, and Bolivia. Wow.

**Steve:** Well, you know, the question I get all the time is, how do I know if my computer has malware or spyware on it? And it'd be neat if you were able to say, oh, just call a nerd, and they'll come over and fix it for you or check it out for you.

**Leo:** They have over 250 core competencies ranging from systems architecture design, I mean, high-end stuff, software development, on-source IT departments, to desktop support. And even, yes, they'll come to your house, residential or small business IT services. Whether it's Cisco, Oracle, PC or Mac, they need it. Fix-it technicians, website designers, programmers, project managers, sales, trainers, security experts – I know there are a few of those listening – antivirus gurus. Just go to Iwanttobeanerd.com. Iwanttobeanerd.com. And by the way, there's a video there of their coffee with you, Steve.

**Steve:** Yeah, I spent, I think, about four hours with them last time I was in Toronto to come up and do the Call For Help show with you, Leo, because I had some time, and I had never met them person-to-person, although we'd had some interaction in email and in teleconferences. So it was really fun.

**Leo:** And there's pictures on there of their cute little red Volkswagen bugs with the Nerds On Site logo in U.K. and Australia and South Africa. I just think it's neat that they're global. This is our first – actually not really because I guess Astaro is also a global company. That's the neat thing about podcasting. We get the global companies. Go to Iwanttobeanerd.com, and you can take the "Are you a Nerd?" test. And we know you'll pass.

**Steve:** Qualify.

**Leo:** And we thank them for supporting Security Now!. So before we get into the meat-and-potatoes of our episode – which is going to be about what, Steve?

**Steve:** We're going to talk, as we promised last week, about the implications and results of Google's really interesting study into their findings about the hard drive reliability within their network of hundreds of thousands of hard drives.

**Leo:** Oh, excellent. Excellent.

**Steve:** Yeah, really neat topic. And, you know, it's not mainstream Internet security or PC security, really. But so many people wrote to me, and even to you, after this Google study in PDF form came out, saying, oh, would you guys please talk about this? It's like, yeah, we can certainly spend a week on this. So that's what we're going to do.

**Leo:** I think it's fascinating. And we should point out, Steve's a security expert, but really I knew him first, well, first I knew you as a designer of the light pen for the Apple II. But after that I knew you and have known you for years as a hard drive expert because of

SpinRite. So this is really your area of expertise.

**Steve:** I will never forget, Leo, that episode of, wait, it was The Screensavers, where you and Kate were standing side by side with a computer. And you weren't sure what she was going to be talking about. But she had some sort of like a Kate's Discoveries segment or something like that. And that's when she discovered ShieldsUP. And so she's talking about this, there's this really cool site done by Steve Gibson. And I think you were sort of looking off-set or something and didn't really hear her. And then you realized she was talking about me. And you said, "Wait, Steve is doing Internet Security? I thought he was hard drives."

**Leo:** Yes, exactly.

**Steve:** It was really fun.

**Leo:** When did that shift happen for you?

**Steve:** It was at a point when the company sort of was going along, SpinRite was doing fine, and we were setting up an ISDN line to our office for the first time to have that level of connectivity rather than all just having modems. And I remember scanning the neighborhood around the IP address we were given, and I found all kinds of C drives, back in the wide-open Windows Filesharing days. And I thought, my God, there's no way people know their C drives are on the Internet. And so I said, okay, I've got to do something about this. And I launched the whole ShieldsUP idea.

**Leo:** That's fantastic. Well, we're going to talk about hard drives in a bit; but first let's cover some errata, some follow-ups. In fact, there's some very interesting stuff to say about our last show.

**Steve:** Well, yes. In fact, we received a piece of email that reminded me that XP has a default setting – remember we were talking about this notion of infrastructure networks versus ad hoc networks.

**Leo:** We were talking about wireless networks. You can have one with a base station, which is infrastructure. But you can also use your laptop or even desktop as a base station, in effect, and that's called an ad hoc network.

**Steve:** Right. In fact, Windows refers to it as computer-to-computer, sort of to explain to people, I think, more what they mean, rather than saying ad hoc. Anyway, the ad hoc networks are a problem. And I had seen several security issues relative to this even before the issue came up last week, like for example people getting infected on airplanes because a bunch of travelers are using their laptops; and, if the laptops are able to connect to each other, there are viruses and types of malware which use ad hoc networks to do cross-computer infiltration.

**Leo:** So the kind of interesting thing is that that connection is the default. It's automatic.

**Steve:** Yes. Thank you, Microsoft. It is the default for XP to connect to either type of network,

either to an access point or to another computer. So what I've done is, on the notes page for this episode, Episode 81 – I'm excited here, Leo, we're heading here towards #100, the big 100 episodes under our belt.

**Leo:** We'll have to have a party.

**Steve:** So on the notes page I have step-by-step instructions showing how you can click on your wireless connector icon down in the Windows tray and step through a couple clicks to get to the dialogue where you're able to say I only want to connect to infrastructure networks, not both ad hoc and infrastructure. And it's really the case that this is wrong. And this is one of those things that XP should really be defaulting to infrastructure-only. But again, in the interest of compatibility and not having people call Microsoft and ask them why they can't connect their computers together, Microsoft has it set so that you can connect to anything, even though most people, I mean, almost all people are connecting to an access point, a so-called "base station," rather than wanting to go computer-to-computer directly without going through a base station.

**Leo:** But let's underscore, it's dangerous to just willy-nilly connect to another computer without asking. And so you do want to disable this. The default is to have it turned on. I want to thank Evan Katz. I don't know how you found out about this, but Evan...

**Steve:** Oh, I think it was Evan, yes.

**Leo:** Yeah, from Manhattan, who sent me the email, and I forwarded it on to you. Actually he sent it to support, sales...

**Steve:** Oh, I got it from every direction, Leo.

**Leo:** Every address he could think of for you. And he talks step by step – in fact, it looks like you're using his screenshots. He gave us some great screenshots. So thank you, Evan, for that tip. Now, what about Vista?

**Steve:** Good question. I dug around in Vista for about half an hour. It's still a frustrating experience for me, Leo, because I'm not using it mainstream myself. I don't expect I will for a while. But I could not find a similar thing in Vista. Maybe Evan will write to us again and tell us where it is in Vista. Although those are my screenshots.

**Leo:** Oh, they're yours.

**Steve:** I didn't see the ones that Evan sent.

**Leo:** Of course they look the same because they're the same dialogue. That's the source of my confusion, all right.

**Steve:** Exactly.

**Leo:** I'll give you credit for that. So we don't know in Vista; and if somebody knows, I'd love to hear about that. And I will dig around a little bit, too, because I have a Wi-Fi laptop on Vista. And I certainly don't want that on.

**Steve:** No. It's something that no one should have on. So it is something that we wish XP had defaulted to off. And then there wouldn't be this bizarre, as we discussed about it last time, this strange Free Wi-Fi...

**Leo:** Free Public Wi-Fi, yeah.

**Steve:** Free Public Wi-Fi that is literally jumping from machine to machine because it ends up being enumerated as available networks, and then it gets stored in the registry, and then that machine will enumerate it to the next one.

**Leo:** It's not a virus, but it is viral in the way it spreads. It's kind of interesting.

**Steve:** There were two other little articles that I picked up in the SANS newsletter this week that I just wanted to share with listeners that I thought would find that sort of interesting, and in some cases distressing. The first is interesting. A man is facing prison – this was picked up by USA Today had the story. It says the subject is "Man Faces Prison for Uploading Movie to the Internet." What happened was, this guy, Salvador Nunez, Jr., his sister is an Oscar screener who received a copy of – I had it right here in front of me, I'm sorry – the movie "Flushed Away," the animated "Flushed Away" movie. Because she got a screener copy, it had a digital watermark in it. She gave it to her brother, who unfortunately thought, oh, let's share this with everybody else and so uploaded it to the Internet. And someone sent a note to the authorities, who found it and tracked it down, were able to backtrack it to her, and he confessed that, yes, he had uploaded it and one other movie to the Internet. So he got caught in a way that was directly trackable due to the fact that there was a digital watermark, a unique digital watermark, in the copy of the movie that he had received from his sister.

**Leo:** They started doing that last year or the year before because they had a real problem with screeners leaking out.

**Steve:** Right. The other story is worrisome. It turns out that law enforcement in Germany – I was thinking about this again when you were talking about how we have a global audience. Law enforcement in Germany is using custom malware to infect the computers of people they want to surveil without their knowledge.

**Leo:** Now, to me, well, certainly in the U.S. that would be illegal. And to me that seems really, really creepy.

**Steve:** Well, it's aggressive. And in fact the story came to light over the – Register.com carried it because a German magistrate denied the German police the right to do this. They apparently went before the court and said, you know, this is what we want to do. And he says, no, this is an unreasonable thing for you to ask. It turns out that there is a division, however, that is still working. It's now called the German Trojan because it's going to be installed in people's machines if the police are able to do this. And apparently there's a bill that is getting ready to be put through that will make this legal, which has sort of resolved this gray area that it's in

right now, and allow the German authorities to install this in people's machines. I don't know, it's not clear to me how they would get it into a machine. That problem still needs to get solved, presuming that it isn't already solved. But it is, it's a creepy development, the idea that there could be software that is known by authorities, installed surreptitiously in people's machines, specially for the purpose of surveilling them.

**Leo:** I can give you a similar example we talked about this week on TWiT. A citizen in British Columbia, 19-year-old kid who wanted to help catch child molesters, hosted a trojan horse posing as a picture of an underage child on several child porn newsgroups, kiddie porn newsgroups, and for the last few years has been monitoring the people who downloaded that trojan. And in fact a judge in your part of the world, Southern California, was successfully prosecuted using evidence that this kid generated by illegally using a trojan horse to get into the judge's computer. Now, I admire the results, but I can't say I admire the means.

**Steve:** Well, in fact, I remember there was a group of security people, among them me, who were talking to the attorney general back in the days of the Internet worms, when the Code Red and Nimda worms were really being a problem. And we were asking whether it was possible to basically fix these machines which had the problem. Since the IPs of the inbound traffic of the worm could not be spoofed, we had the IPs of all the machines that were infected. So it was so tempting to create our own inoculation, essentially, for those machines, and go and remove this software, fix this problem in the machines that were infected. And, you know, anybody running a honeypot could have collected the IPs of the infected worm-transmitting machines. Anyway, of course the word came back in no uncertain terms that it was absolutely illegal for us to make any unauthorized modification to a machine over which we don't have any rights at all, even when we're doing something good, something that the owner would almost certainly want us to do, like fix their machine for them because it's actively broadcasting Internet worms. It's like, no, you know, there's no way that's okay.

**Leo:** Yeah. Very interesting. Very interesting stuff. And certainly an ethical issue that we haven't heard the last of. Shall we get to hard drives?

**Steve:** Sure.

**Leo:** Any other errata?

**Steve:** No.

**Leo:** Before we do, I just want to mention that I had a very good experience with a certain program this week known as SpinRite.

**Steve:** That is so cool, Leo.

**Leo:** Yeah, well, I'm just glad I know you, that's all I can say. Because I forgot my – I lost my serial number. And I had burned a CD, and this is one of the good things you can do with SpinRite is burn a CD because you do want to boot to the SpinRite CD so that you can be in DOS when you're doing the checking. And I don't know what I did with my CD, and I couldn't find it. You know, and I realized later I shouldn't have bugged you. I could have

gone to the website and retrieved my serial number that way. Thank you for being very patient with me. But the good news is, the computer that was dying is the one I'm using right now to record this, I record all the podcasts on. It's a pretty important computer. And it wasn't booting. And it said there was something, you know, couldn't find an operating system. So I ran SpinRite on it. I ran the long maintenance, you know, the check and the maintenance fix overnight on all the drives. It's been flawless ever since.

**Steve:** I just love that, Leo, that is so cool, so cool.

**Leo:** I am grateful to you. Well, I just thank you so much.

**Steve:** Well, and for what it's worth, the ecommerce system that I wrote, of course, in Assembly language to run all of this, if any user holds onto their transaction code that they receive when they purchase SpinRite, that enables them at any time in the future to go back and grab a copy of SpinRite. So if they were somewhere else, I mean, you just use a web browser. So you could even do it if you were away from home and your laptop started giving you trouble. And also for people who don't have that written down in their wallet or have lost it, they can certainly contact our sales email at any time, give us enough information to find you in our database, and we'll just let you know what your transaction code is, and then you're off and away. So even people with less direct connection to me, Leo, than you have are similarly able to get themselves back to their copy of SpinRite if they don't have it when they need it.

**Leo:** We all have a friend at SpinRite. GRC.com is the website. If you want to read testimonials, SpinRite.info. I can't recommend it more highly. The ultimate disk recovery and maintenance utility. And I've used it before to save drives. This is just the most recent example. It's such a great thing to have in your toolkit. So, Steve, speaking of drives...

**Steve:** While we're on the topic...

**Leo:** While we're on the topic. Now, we should say this isn't exactly a security topic. But I think when you have kind of one of the foremost experts on hard drives available, that it's important, it behooves us to talk about this Google study. Let me just ask right upfront, when you read this, were you surprised by the results?

**Steve:** Not one bit. One of our favorite expressions at GRC for years has been that "SMART is dumb."

**Leo:** I like that. You're talking about the SMART drive capability built into hard drives.

**Steve:** Yes, the Self Monitoring and Reporting Technology, SMART. Or I'm sorry, Smart Monitoring Analysis and Reporting Technology. The history of this so-called SMART technology is not well known. Back in the early days of the PC, when drives were beginning to become sort of, I want to say self-aware, you know, like Skynet with "Terminator." But when drives were beginning to get intelligent, remember that the earliest drives you actually had a disk controller with a whole bunch of stuff on it, and then one or two cables that went to the drive, sometimes a serial cable and a radial cable. So you'd daisy chain one, and then you'd have a direct radial cable from the controller to each of your drives. And these drives were really – they were dumb drives. They were spinning platters, read-write electronics, and in many cases a stepping motor

that just stepped the heads in and out. And those drives held all of 10 megabytes or 20 megabytes or, you know, 40, even, of data.

**Leo:** That was a big day. I remember when they went from 20 to 40.

**Steve:** It's like, wow, yeah, now we've got...

**Leo:** [Indiscernible] RLL, and that was a big jump, oh.

**Steve:** Right. Right. So what happened was, when drives were dumb like that, there wasn't anything you could really do on the drive. When drives went to the IDE interface standard, what happened was essentially, well, in fact, IDE stands for Integrated Drive Electronics. The controller was moved onto the drive, that is, like a disk controller was moved onto the drive so that all of the data flow, the serial data flow on and off the heads and the control, that was all made local. And an interface was then created that allowed you to use a simple parallel cable down to the motherboard, or at that time to even a very simple IDE controller. But the idea was the drive started being smart, and had its own microprocessor on it. And so now it was sort of autonomous.

Well, Compaq was the big leader, as you remember, in the clone market at the time. They were the first people who came up with a clean room-created BIOS, that is, a BIOS clone of the original IBM PC that had never – where the authors of the BIOS had never had any contact with the source code which IBM published in their technical reference manual. So in order to avoid any possibility of copyright infringement, Compaq put a bunch of engineers, literally locked them up in a room, gave them – and these are people who had never even seen the source code for the IBM BIOS. But from the BIOS's specification, they wrote their own BIOS from scratch.

So Compaq, of course, took off as a major clone of the IBM PC. IDE drives began happening, and they began failing. And what Compaq did was essentially they turned around to the drive industry and said, look. I think this was Seagate, Quantum, and Connor Peripherals at the time were the three, and Western Digital joined in later. Compaq said, look, we're going to tell you what you have to do. Well, the drive manufacturers were not excited about being told what they were going to do by Compaq. But Compaq's purchasing power was king at the time. I mean, Compaq was buying so many drives that the manufacturers said, well, okay, what do you want? And Compaq said, we need some way of knowing what's going on inside the drive because, as the drives became intelligent, that allowed them to hide what was going on inside. And so Compaq felt like, wait a minute, we don't know what's happening now because there's this intelligence in the drive that has created sort of an abstraction of a hard drive storage. And of course that's even more true these days, when drives are, like, self-relocating sectors and doing all kinds of stuff autonomously. They really do present this opaque wall. So Compaq said, you need to give us a way to know what's going on behind the scenes in the drive.

Well, so this SMART specification was created to allow a means for the drive to publish some information about itself. The problem is that, in order to get the manufacturers to agree, and they were not happy about this, the specification had to be left very loose. You know, it would be like, well, we'd like you to tell us, sort of give us some health parameters that, when they go lower, that's worse, and when they go higher, that's better. And Compaq tried to nail the manufacturers down to a tighter specification, but it just wouldn't work. So what we've ended up with, even today, 15 years later, is a specification which is sorrowfully weak. And in fact, the Google paper talks about this in several instances where they are seeing no uniform meaning behind many of these SMART parameters across a large install base of hard drives.

**Leo:** Well, they even go farther. They say a lot of the failed drives had no SMART errors at all.

**Steve:** Well, now, that's sort of a different issue. That's the issue of how effective is all of this, as opposed to what is...

**Leo:** Is it predictive, yeah.

**Steve:** What is all of this, yeah. And of course that certainly – that's a perfect segue, Leo, because certainly the goal of SMART was to allow software running in the machine to keep an eye on what the drive was doing and predict hard drive failures. That was the whole idea. The problem is, and in fact the reason that we've always used the slogan "SMART is dumb" at GRC, is we see dead drives all the time where the SMART system says, oh, everything's fine.

**Leo:** It's all working, no problem.

**Steve:** Exactly. The drive can die spontaneously while SMART is completely happy and sees nothing going wrong. At the same time there are drives which look like they're on their last throes from a standpoint of the SMART data that just keep on going for years. So the problem really, and this is what Google's findings were – well, actually Google found a lot of things. But relative to SMART they found a very disappointing lack of correlation between, well, based on extensive statistical analysis of over 100,000 drives within their system, a very weak analysis...

**Leo:** Correlation.

**Steve:** Thank you, yes. A correlation of what SMART showed to the actual failure of the drives. The reason I was pausing there is I was trying to find the exact reference. It said in the study, it says – and I'm reading from their report on the predictive power of SMART parameters. It says, "Given how strongly correlated some SMART parameters were found to be with higher failure rates, we were hopeful that accurate predictive failure models based on SMART signals could be created. Predictive models are very useful in that they can reduce service disruption due to failed components and allow for the more efficient scheduled maintenance process to replace the less efficient (reactive) repair procedures. In fact, one of the main motivations for SMART was to provide enough insight into disk drive behavior to enable such models to be built." Which is a long-winded way of saying we hoped that SMART would allow us to know when a drive was going to die before it does.

**Leo:** Well, they weren't alone. I think we all hoped that.

**Steve:** Yeah.

**Leo:** But I think we all knew, anybody who had any experience with drives, and I've said this for a long time, knew that really SMART didn't seem to have that kind of ability.

**Steve:** And if it did, we'd all be using it a lot more than we are. Because it's there, but it's not highly used. And just to finish this one thought here from the report, reading from the report, it

says, “After our initial attempts to derive such models yield relatively unimpressive results, we turn to the question of what might be the upper bound of the accuracy of any model based solely on SMART parameters. Our results are surprising, if not somewhat disappointing. Out of all failed drives, over 50 percent of them have no count in any of the four strongest SMART signals, namely scan errors, relocation count, offline relocation, and probational count. In other words, models based only on those signals can never predict more than half of the failed drives.”

So essentially what Google found is that many drives were failing, more than half of theirs were failing where nothing showed up at all in the SMART subsystem; and also that exactly the reverse was happening, is that SMART was showing things where drives never failed. There were things they found, for example, when they would ask the SMART system to scan the drive, if an error was found during that scanning, the drive was 39 times more likely to fail in the next 60 days than all other drives. Except that it turns out 39 times more likely wasn’t predictive enough to say, okay, we should replace the drive, because it turns out that there were lots of drives that had scan errors that never failed. So, I mean, it just...

**Leo:** It’s useless.

**Steve:** It’s basically, yes. It’s funny, too, because in SpinRite 6 I incorporated, as you probably saw when you were running SpinRite just recently, Leo, I incorporated a real-time monitor of the whole SMART subsystem. Now, there are interesting things that you can see. SpinRite will show you these SMART parameters being driven down. One of the things that makes SpinRite different than just scanning passively is it’s reading and writing and reading and writing and writing to the drive, essentially, you know, really working the surface. So rather than just being a read-only scan, SpinRite is a full read-write stress that inverts all the bits, writes them back, reads them back, reinverts them, writes them back, reads them back. So it’s basically checking the entire surface aggressively. So what’s really interesting is that we’ll often see that the SMART parameters which are normally running okay on weak drives, sort of these health parameters get driven down for a while while SpinRite is running on the drive, and then gradually recover. If they don’t go critical, it’s probably not something to worry about. But if you were watching it happen, you could compare this behavior of your drive to its behavior, for example, six months from now. And if it started seeming like it was worse, then you’d get some qualitative sense that the drive was, you know, not doing very well.

So anyway, unfortunately, just the SMART stuff by itself, as Google found and as all of our experience has shown, just isn’t useful enough. And imagine software saying, oops, you need to replace your drive. Well, who’s going to believe that? Certainly in a mission-critical situation where you absolutely can’t have a drive fail, maybe that’s interesting. Except that what Google found is also that there was some infant mortality behavior, that is, they found that younger drives just installed had a higher failure rate for the first three and then even six months than drives that had been there for a year. So the idea was that newly minted drives could have something wrong with them that has sort of gotten through quality control and manufacturing that would manifest to the detriment of the drive during its first few months of use. After that time you get to a point where, okay, the drive is now sort of matured, and it’s going to give us a nice long life.

What designers have always assumed was that the statistical curve would look – it’s called a “bathtub curve,” where it’s high initially. It drops down after you move out of the infant mortality period. Then maybe hopefully for five or six years you have relatively lower failure rate. And then at some point later on, just due to the drive’s aging and fatigue and mechanical wear and tear, the reliability begins falling again, and the rate of failure goes back up. Thus the other side of this so-called “bathtub curve.” What Google found, and this was what we briefly discussed last week, and it’s what’s so disturbing, is that in fact drives two years old were experiencing suddenly much higher failure rates than drives one year old, and three years old were even more. You may remember, I think it was 8.6 percent annual failure rate, meaning

two out of every 25 drives is failing every year at three years old.

**Leo:** Well, it's kind of still a bathtub curve, it's just a shallow – it's a narrower bathtub than we would expect.

**Steve:** It's a one-year bathtub instead of, like, five or six years, exactly.

**Leo:** Is this the biggest study ever done of hard drives? It must be.

**Steve:** Well, yes. There had been lots of littler studies. Actually filed during the same conference, also in the 5th USENIX conference, was another paper, and we've got a PDF link to it also on our notes page for Episode 81. Two guys, researchers at Carnegie Mellon University, CMU, studied very much the same sort of data. They didn't use their own 100,000-drive dataset, but they got maintenance data from a number of very large ISPs that were willing to make this available. And so they did a similar study, and their conclusions were very much the same. One of the things they found was that they were talking about how manufacturers would quote a one million hours mean time to fail, MTBF or MTTF, Mean Time to Fail or Mean Time Before Failure. Well, a million hours equates to 0.88 percent annual failure rate. You may remember that that Seagate spec I referred to, which surprised me, it was for a Barracuda 7200.9 drive that quoted a 0.34 percent annualized failure rate.

So anyway, what the CMU guys found was almost the same as what Google found. These guys were not looking at SMART data because they weren't doing real-time SMART capture in their whole infrastructure the way Google was. They were just looking at, like, basically aggregating, did a big statistical study of maintenance reports. But reading from their conclusions, they say, "Large-scale installation field usage appears to differ widely from nominal data sheet MTTF conditions. The field replacement rates of systems were significantly larger than we expected based on data sheet MTTFs, which was..."

**Leo:** What a surprise.

**Steve:** Yeah. Which sort of is to say, oh, you mean manufacturers are not telling us...

**Leo:** Inflating the numbers.

**Steve:** ...what we should have – they also say...

**Leo:** Instead of .88 percent, it was as much as 13 percent. Well, that's close.

**Steve:** Yes. I mean, substantially higher. They say for drives less than five years old, field replacement rates were larger than what the data sheet MTTF suggested by a factor of two to ten.

**Leo:** Wow.

**Steve:** And for five- to eight-year-old drives, field replacement rates were a factor of 30 higher than what the...

**Leo:** Unbelievable.

**Steve:** ...than what the datasheet MTTF suggested. So basically what, I mean, this is sort of corresponding, unfortunately, with people's real-world experience. And, I mean, you just had a drive that wouldn't boot on you.

**Leo:** Right.

**Steve:** Every week or so I share a real-life experience from one of our customers who turns the computer on, and it says, I'm not a computer anymore. Operating system not found or whatever.

**Leo:** That's exactly what happened to me.

**Steve:** And I think I read the number somewhere, I think it was 350 million drives were manufactured in 2006. So we have, obviously, there's hard drives all over the place. And coming back to your point about how this is not really about security, I guess my feeling is, well, certainly that's the case. But there's almost nothing more important to computer users than their hard drives.

**Leo:** When in fact that's the number one cause of data loss. It's not hackers. It's dying hard drives.

**Steve:** Exactly.

**Leo:** Can we infer anything else or learn anything else from these studies? I know that Google looked at temperature, for instance.

**Steve:** Yes. And that was really interesting, too. My wisdom has always been that temperature matters. I've seen drives failing due to high temperature conditions.

**Leo:** Now, by the way, I think that's what was wrong with my drive. It's a shuttle case, and it's a very small, tight case.

**Steve:** Ah, yes.

**Leo:** And almost immediately SpinRite said, hey, you've exceeded 150 degrees. I'm going to stop.

**Steve:** Yes, in fact, it's really interesting. That's one of the things that laptop users often report

with SpinRite. I actually have a photo that was sent to me a couple weeks ago of a laptop running SpinRite in a refrigerator. Somebody took out...

**Leo:** Not recommended, by the way.

**Steve:** They took out all their food and a couple shelves, and they stuck the laptop in the refrigerator. The problem is, laptops are notoriously troubled about cooling. I mean, I'm amazed, when you put a laptop on your lap, literally how hot these things get.

**Leo:** Well, I bought a 7200 rpm drive in my most recent Dell, and that thing, before it even starts working, really, is already blowing very hot air out the side, you know, the fans are working hard. And I'm sure the hard drive has a lot to do with it.

**Steve:** Yes, it certainly does. And so what happens is it turns out it's the seeking of the drive. When a drive's head actuator is not moving, the only real power consumption you have is the relatively steady state spin of the drive. But from an engineering standpoint, or from a physics standpoint, when you need to move the head very quickly to another cylinder, you need to mechanically accelerate the head up to speed, it flies across the cylinder, and then you need to stop it immediately. So what you're doing is you are briefly applying very high forces laterally to the head in order to move it from a rest state into motion and back. So anything, even like defragging, for example, which is putting the drive through a great deal of exercise, will increase the power consumption of the drive, but also a lot of heat. The energy that is dumped into the drive comes out acoustically; thus you hear, you know, you actually hear the drive doing things. And thermally is the other way that the energy is converted from electrical energy. So you end up seeing that the drive's temperature will increase. So SpinRite, using the SMART system, is constantly keeping an eye on the drive temperature to alert people that, you know, hey, you've got a problem here with your drive getting beyond manufacturer's spec. And when SpinRite says it's a problem, it really is. I mean, I set that to the upper limit of manufacturers' safe running conditions for a drive. So if anyone can do – if SpinRite says your drive is running too hot, anything you can do to move more air across it or keep it cool, just for normal daily use, too, would be good.

**Leo:** In fact, that's what I did is I opened the case up, and I have a fan flowing air from the window across the whole system.

**Steve:** Now, here's the problem with Google's study, which they acknowledge, and with, for example, quote, my common wisdom about temperature. And that is, their study was based on many years of aggregate data collecting, during which time new drives came along, new makes and models, and manufacturer mixes changed. I mean, the most potentially interesting thing Google could have told us and deliberately didn't, they referred to it as "proprietary information," was to show us a breakdown of drive failures by manufacturer.

**Leo:** Oh, I would have loved to have known that, yeah.

**Steve:** Who would not kill to have that information. I have my secret suspicions of who would have been top of the list for drive failures. But anyway...

**Leo:** Well, who?

**Steve:** Well...

**Leo:** You've got a database. I mean, you know.

**Steve:** Yeah, I never buy Western Digital drives.

**Leo:** Oh, that's funny, because I do.

**Steve:** And the reason I hedged, and I didn't offer that initially, Leo, is that I have said that before, and I have had people say, wait a minute, I've never had a WD drive give me trouble.

**Leo:** No, in fact, I put Western Digital Raptors in every single new computer. In fact, our web server is running on Raptors in RAID-5. I've got three Raptors on each web server.

**Steve:** Well, and the Raptor is a high-end WD drive.

**Leo:** You pay a lot more for it, yeah.

**Steve:** Yes. I do think that, although the Carnegie Mellon study, the CMU study, did not show that SCSI drives were less error-prone that SATA drives or than Fibre Channel drives. So they did take a look at, by interface type, what drives tended to fail more.

**Leo:** But that makes sense. You wouldn't expect the interface to impact that, would you?

**Steve:** Although my sense, again, has been that SCSI drives were higher end, more expensive, and just they were built with a little more love and care than push-them-out-through-the-consumer-retail-channel IDE and SATA drives. But here's my point. My feeling is that, and certainly this is the case in the last five to seven years, drive technology is clearly a moving target. These drives are changing internally to such a degree that I'm not really sure that a study taken over a long period of years really tells us anything about today's drives. Because, for example, the Google study did show that older drives showed more of a temperature susceptibility, that is, more of a problem to temperature than newer drives. So you might assume, then, first of all, that okay, maybe as any drives get older they become more temperature susceptible, though I don't know why that would be the case. Or you could assume that newer drive technologies are always going to be less susceptible to temperature than older drive technologies were. So, for example, five years from now a repeat of this study on a next-generation population base of drives would yield different results.

**Leo:** They said that was actually the greatest, the most surprising result of the study was that temperature didn't seem to have much to do with failure rate.

**Steve:** Yes. And, well, and what's bizarre is, it turns out that the failure rates were higher for cooler drives. There was actually – you had to get the drives very hot before the reliability finally spiked in the negative, or running drives very cold. They just didn't seem to like that. Which of course is not much of a problem because they heat up pretty quickly.

**Leo:** Right, right. A very interesting study. And as Steve mentioned, there are copies of this study and the CMU study from Bianca Schroeder and Garth Gibson on his website. So you can read this. This is all from FAST '07, the 5th USENIX conference on File and Storage Technologies, which just happened over the Valentine's Day weekend.

**Steve:** Right. So my takeaway from all of this, I mean, it would be very cool if we could tell people before their drive was going to fail, with a high degree of reliability, that it was going to. On the other hand, if you cry wolf once or twice, no one's ever going to believe you again. And the common wisdom is, of course, back up your data. And so where I am with SpinRite, and where we've always been, basically, is saying to people, look, the best thing you can do is keep a current backup. That's really your only way of knowing that your data's going to be safe. But people never get around to it, or their backup's a few months old, and they did some really important stuff in the last few months that they have to have back. So the good news is, with SpinRite, there's a way of at least turning back that clock a little bit, maybe getting, you know, if it doesn't completely fix your drive, it'll get you back going again long enough to pull your critical data off. So...

**Leo:** Well, you remember what HAL 9000 said in "2001": "I would recommend that we put the unit back in operation and let it fail. I would recommend...." So even then, even "2001" that was a recommendation, so that's pretty much the way it is. And certainly you're not going to pull a drive just because SMART says to. That would be a big mistake. Main thing is back up, for crying out loud. Just make backups.

**Steve:** Right. And more often than not, SMART won't tell you that a drive is failing, and it does anyway.

**Leo:** Well, we thank you for summarizing. I think it's fascinating stuff. And I hope you folks don't mind that we diverted a little bit from our security topic; although, as Steve pointed out, there is nothing more important than keeping your data secure. And certainly your hard drive is the front line of defense there.

We are brought to you, as always, by the good folks at Astaro, the makers of the Astaro Security Gateway. They just announced Version 7, some significant new improvements. We've talked about them before, things like transparent email encryption and decryption, which means your entire enterprise can be using email encryption without having to worry about teaching everybody how to do it individually. It's just fantastic. Also you'll find that signatures are enabled using S/MIME or OpenPGP standards. I use S/MIME myself, it's great. Of course I have an Astaro Gateway, that's why. Inbound mail is also automatically decrypted. They've got remote access via SSL; that's significant. You've got IPSec, L2TP over IPSec, PPTP tunneling with SSL, makes it so much easier to enable. And of course all sorts of useful things, including scalability via clustering; the usual antispam, firewall, intrusion detection, anti-phishing; I mean, it's just a really superb unit. And you can try it absolutely free by calling Astaro or visiting them online at Astaro.com, or call 877-4AS-TARO to schedule a free trial of the Astaro Security Gateway, now in Version 7, in your business.

By the way, you can now – and this is, I think, really fantastic news – download, if you're a noncommercial user, not only – you could always download the ASG software because it's open source and free for noncommercial use. But now they're going to include all subscriptions and Astaro Up2Date. This used to be a 79 euro per year subscription. Now that's absolutely free. Just another reason to visit Astaro.com. We thank them so much for their support of Security Now!.

So I'm going to keep my Western Digital Raptors, Steve Gibson, I'm sorry.

**Steve:** That's okay. Well, and my favorite drives were always, when I decided that – because I had several WD drives die on me. And I thought, okay, when am I going to learn a lesson? I switched over to Quantum. And I've loved Quantum drives for years. Then Maxtor bought Quantum. And then, as you know probably, recently Seagate bought Maxtor.

**Leo:** Right, so it's all the same now. There's really only two companies, basically. I guess Hitachi and IBM are still in business.

**Steve:** Well, and what, I think Hitachi bought IBM's drive business.

**Leo:** Oh, did they? Oh, all right.

**Steve:** Yeah. So the problem is, drives are phenomenally complex. There's so much to go wrong with them that, you know, they fail from time to time.

**Leo:** And I should mention that the drive that I had problems with was in fact not the Western Digital Raptor. I think it was a Maxtor drive. The Western Digital's been working in there flawlessly. And I have one in my Mac, one in my PC. Every machine I just buy a Raptor because I like – now, tell me if I'm wrong, but I believe – this could be a completely fallacious belief – that the 10,000 rpms makes a difference. It also has a 16-meg cache in speed; right?

**Steve:** Oh, in speed, absolutely. You will definitely be getting a transfer rate performance because literally you're spinning the drive faster. So not only do you have a higher data rate, but you have a lower latency because it's going to take less time for the sector transfer to begin.

**Leo:** Yeah. That's why it's not my boot drive, it's my media editing drive. We record to the Raptor, and we edit on the Raptor, because I want all the speed I can get.

**Steve:** Well, again, I didn't mean to badmouth WD. I know there are other people who really like Western Digital drives. So it's just like, to each their own. It's like people have cars they like and cars they don't like, so.

**Leo:** I think also that it's probably true across the board, the cheaper drives are going to be less reliable. Higher density cheaper drives.

**Steve:** Oh, one thing I did pick up when I was doing the research for today's show – although obviously I live and breathe this stuff, so I knew a lot of it. But one of those studies talked about the lower reliability for higher headcount. And so there is the feeling I had had intuitively, that smaller drives were more reliable.

**Leo:** But we thought it might be areal density. It turns out to be the heads that matter.

**Steve:** Well, yes, because the smaller drives only have a single platter. And what they're doing is they're adding platters. Which is why drives jump typically, for example, by 80 gigs or in some cases 40 gigs. It'll go 40, 80, 120, 160 because they're adding surfaces to the drive. And it turns out, of course, that the more of that stuff you've got in there, any one of them failing will bring the drive down. So the fewer you have, the more reliable it is.

**Leo:** Excellent. This stuff is good to know. That's why we always listen to Steve Gibson. If you want copies of the studies and more information, including information on how to turn off ad hoc networking...

**Steve:** And anybody using Wi-Fi with XP absolutely should do this. You want to go to our notes page and follow that.

**Leo:** That's at GRC.com. That's where you'll also find 16KB versions of this show for the bandwidth-challenged, and Elaine's fantastic transcriptions. By the way, Steve, you've inspired me. We're going to do a deal with a company to do transcripts of many of our other shows, too, because I think people really do enjoy reading along often with the show.

**Steve:** I get a lot of great feedback about it. And of course it makes it also text-searchable. So Google and other search engines are able to find those articles and refer people to them.

**Leo:** Yeah, I think it's a good thing. GRC.com. That's where you'll also find, of course, Steve's free security programs like SecurAble and ShieldsUP, and his day job, SpinRite, a great recovery and maintenance utility.