IN THE UNITED STATES DISTRICT COURT
DISTRICT OF NEVADA

THE SCO GROUP, INC.,
A DELAWARE CORPORATION,

Plaintiff,

VS.

Civil Action File No.
CV-S-04-0237-RCJ-LJL

AUTOZONE, INC., A NEVADA
CORPORATION,

Defendant.

VIDEOTAPED DEPOSITION OF BOB CELMER

MAY 9, 2005

ALPHA REPORTING CORPORATION
KORIAN NEAL, RPR, CCR
100 North Main Building, The Lobby
Memphis, Tennessee 38103
(901) 523-8974

The videotaped deposition of BOB CELMER, taken on behalf of the Plaintiff, pursuant to Notice, on May 9, 2005, beginning at approximately 9:00 a.m., in the law offices of Baker, Donelson, Bearman & Caldwell, 165 Madison Avenue, Suite 2000, Memphis, Tennessee.

This deposition is taken in accordance with the terms and provisions of the Federal Rules of Civil Procedure.

All forms and formalities are waived, and objections alone as to relevancy, materiality and competency are reserved, to be presented at or before the hearing. Objections as to the form of the question must be made at the time of the taking of the deposition. The signature of the witness is not waived.

- APPEARANCES -

For Plaintiff: DAVID STONE, ESQ.
CHRIS IANNICELLI, ESQ.
Boies, Schiller & Flexner, LLP
Attorneys at Law
150 JFK Parkway
Suite 100
Short Hills, New Jersey 07078
(973) 218-1111

KEVIN McBRIDE, ESQ.
1299 Ocean Avenue
Suite 900
Santa Monica, California 90401
(310) 393-1080

RYAN TIBBITTS, ESQ.
SCO General Counsel
355 South 520 West
Suite 100
Lindon, Utah 84042
(801) 765-4999

For Defendant: DOUGLAS J. STEWART, ESQ.
DOUGLAS BRIDGES, ESQ.
Alston & Bird, LLP
Attorneys at Law
One Atlantic Center
1201 West Peachtree Street
Atlanta, Georgia 30309-3424
(404) 881-7000

ALSO PRESENT: REBECCA W. BALLOU, ESQ.
JOHN DUBOIS
WILL SMITH, VIDEO SPECIALIST

COURT REPORTING FIRM: ALPHA REPORTING CORPORATION
Korian Neal, RPR, CCR
100 North Main Street
The Lobby
Memphis, Tennessee 38103
(901) 523-8974

INDEX

Witness (BOB CELMER) Page
Direct Examination by Mr. Stone.....06

EXHIBITS

No. Description Page
32 Notice to Take Deposition.....07
33 Response to Interrogatories.....32
34 E-Mail.....48
35 Response to Supplemental Interrogatory..56
36 Supplemental Information.....68
37 USB Hard Drive Information.....106
38 Groklaw Posting.....153

COURT REPORTER'S CERTIFICATE.....163

VIDEO SPECIALIST: Today is May 9th, 2005, 9:05. Location is Baker, Donelson, Bearman & Caldwell, 165 Madison Avenue in Memphis, Tennessee. My name is Will Smith, certified video specialist with Alpha Legal Productions, located in the lobby of the 100 North Main Building in Memphis.

This case is entitled The SCO Group, Inc. versus AutoZone, Inc. Our deponent today is Mr. Bob Celmer. This video deposition is requested by the plaintiff's counsel, Mr. David Stone. Will counsel and all present please identify themselves for the record at this time.

MR. STONE: David Stone and Chris Iannicelli from Boies, Schiller & Flexner for the plaintiff, SCO Group.

MR. McBRIDE: Kevin McBride for the plaintiff.

MR. TIBBITTS: Ryan Tibbitts, general counsel, and John DuBois, both from the SCO Group.

MR. STEWART: David Stewart and Doug Bridges, Alston & Bird for the defendant, AutoZone.

MS. BALLOU: Rebecca Ballou from AutoZone.

MR. CELMER: Bob Celmer, senior technology advisor for AutoZone.

VIDEO SPECIALIST: The deponent may now

be sworn in by Ms. Korian Neal of Alpha Reporting Corporation.

BOB CELMER,

having been first duly sworn, was examined and testified as follows:

DIRECT EXAMINATION

BY MR. STONE:

Q. Good morning, Mr. Celmer. How are you?

A. Good. How are you?

Q. I don't need to tell you my name because I already told you my name. So we'll cross that off my list. You're here today for a deposition. Have you ever been deposed before?

A. Yes, I have.

Q. Okay. So you're familiar with how depositions proceed, that I'm going to ask you questions, you're going to answer the questions. From time to time, your counsel may object, and you still must answer the question unless he instructs you not to answer the question. You understand that?

A. Right.

Q. You also understand that you're here under all the oaths of pains and penalties of perjury that would apply if you were testifying in court.

A. Right.

(WHEREUPON, A DOCUMENT WAS MARKED AS EXHIBIT NO. 32, AND IS HERETO ATTACHED.)

Q. Mr. Celmer, I'm going to show you what's been marked as Celmer Exhibit 32, which is a letter from myself to Mr. Stewart along with a notice for a 30(b)(6) deposition, ask you if you've ever seen that before today.

A. Yes.

Q. And you had an opportunity to review it before you are appearing here today?

A. Yes.

MR. STONE: And I understand, Mr. Stewart, that the defendant has filed certain objections to this notice, and I'm not meaning to ignore those objections by not referring to them or marking them. I just don't happen to have them here. But you know, I understand you have your objections, and you'll raise them at the appropriate time.

MR. STEWART: That's fine.

Q. (BY MR. STONE) Mr. Celmer, the purpose of this deposition is to find out facts pertaining to SCO code from an operating system called OpenServer. Are you familiar with that operating system?

A. Yes, I am.

Q. And you're familiar with the fact that AutoZone migrated or moved from using that system to using a type of a Linux operating system at some point in the past?

A. Yes.

Q. And currently, is AutoZone using Linux systems on all of its store servers?

A. Not all of its store servers.

Q. So some of the store servers use other types of operating systems?

A. Yes.

Q. And were you involved in any way in that migration process?

A. Yes.

Q. Okay. Well, why don't we start out by you giving me a little bit of your history. Start with your education. Where did you go to college?

A. I went to Southern Illinois University as an undergraduate and was awarded my master's degree from Memphis State University here in town.

Q. When did you graduate from Southern Illinois?

A. 1984.

Q. And when did you achieve your master's?

A. 1990.

Q. And what was your master's in?

A. Computer science.

Q. Did you study -- when you were studying for your master's, did you study the Linux operating system?

A. No, we didn't.

Q. Did you study the OpenServer operating system?

A. No.

Q. Did you study Unix-based operating systems?

A. Yes.

Q. When you -- well, what is the first job that you held?

A. First job I held was a lot boy on a used car lot.

Q. Okay. And when --

MR. STEWART: Do you mean first job after college?

Q. (BY MR. STONE) What was the first job after college that you held?

A. Construction laborer.

Q. Okay. And for how long did you have that job?

A. Nine or ten months, I suppose.

Q. And what's the next job that you held?

A. Security guard supervisor.

Q. Okay. And how long did you have that job?

A. Four years.

Q. So tell me what year we're in now.

A. 1989 --
Q. Okay.
A. -- I believe is when I left that job.
Q. And what job did you hold after that job?
A. Let's see. I was a programmer for a company whose name I can't recall right now. It was a company here in town that provided software for credit unions. I'm sorry. I'm drawing a blank on the name of the company right now.
Q. And so you programmed software that assisted credit unions?
A. Right.
Q. How long did you hold that job?
A. About -- less than a -- less than a year. About six to nine months.
Q. And where did you go next?
A. Malone & Hyde, a grocery distributor.
Q. And you programmed software for them?
A. Right.
Q. Where did you go next?
A. AutoZone.
Q. And what year was that?
A. December of 1990.
Q. So the first time you joined AutoZone was in December of 1990?

A. Right.
Q. And what was your position at that time?
A. Programmer.
Q. And did you -- what were your responsibilities as a programmer?
A. The development and maintenance of programs for our store systems.
Q. And at that time, was AutoZone using OpenServer systems?
A. No. At that time, we were using Zenix.
Q. And was Zenix a Unix-based operating system?
A. I believe Zenix is a Unix-based operating system.
Q. Do you know who you were licensing Zenix from?
A. The Santa Cruz operation.
Q. And what's the next job you held after programmer?
A. My first title change after programmer I believe was project leader.
Q. And when was that, approximately?
A. 1993, probably.
Q. And how did your responsibilities change as a project leader from a programmer?
A. The only real change in responsibility was the review of junior personnel in an official capacity.

Q. So you took on certain management responsibilities that you didn't previously have? Is that correct?
A. That's right.
Q. Did you receive any further promotions after that time?
A. Yes.
Q. When's the next promotion you received?
A. Probably 1998 or so, I was given the title of technical fellow.
Q. And technical fellow, what were your responsibilities as a technical fellow?
A. I really had two responsibilities as technical fellow. One is, it was an experimental position. We wanted to find out whether it was appropriate to create a separate technical track for the developers at AutoZone, technical career track. And so a large part of my responsibility was to outline duties and responsibilities at the upper end of that track.
Other than that, I was involved with Y2K preparations, and I was investigating a multi-dimensional database product for use in data mining.
Q. Okay. And did you receive any further promotions after that date?

A. No. My current title is sort of a standardization of that technical fellow title.
Q. Okay. When is the first time that you recall that AutoZone began using any version of the OpenServer operating system?
A. I really don't recall the year. When we migrated from Zenix to Unix, it wasn't called OpenServer. And then with the subsequent -- subsequent release of the operating system, SCO changed the name, I believe, to OpenServer.
Q. Okay. So for some period of time, you were using a new operating system that was different than Zenix, and then at some point, the name of that operating system was changed to OpenServer. Is that correct?
A. Right.
Q. Do you recall when that happened?
A. No. I'm sorry, I don't.
Q. Was it some time between 1995 and 2000?
A. I really don't remember when it was.
Q. Okay. And then at some point, AutoZone decided that it was going to migrate again or change again the operating system it was using to Linux. Is that correct?
A. We wanted to explore -- to investigate Linux

to see whether it was sufficiently stable and reliable for our use. And one of the tests that we wanted to make of Linux was to see whether it would support at least our most critical applications and to find out how difficult it was to port the software.

Q. And when did you first begin those inquiries about Linux?

A. I think we began to port the software around 1999.

Q. And -- but before you began to port the software, you did some investigation, is that correct, into Linux?

A. Yes.

Q. And approximately when did you begin your investigation, if you recall?

A. Very fairly early on, we had a copy of Linux, early 90s, '93, '94, I believe.

Q. So you believe you were at least looking at the possibility of an alternative of Linux as an operating system as early as '94, '95?

A. I think so. Yes.

Q. Were you actually involved in the decision to migrate AutoZone's machines to the Linux operating system?

A. Yes.

Q. Who made that decision?

A. Our senior management at AutoZone would have had to make the decision. So that would have been one of the vice presidents at that time.

Q. Do you recall who the person -- the name of the person?

A. In general for critical decisions, we tend to act as a team. And so our CIO at that time clearly would have, you know, been on the hook for the ultimate decision. That would have been -- we had a CIO transition around that time. I don't remember whether Bruce Clark was the CIO or whether Steve Valentine was the CIO. But in either case, they would have had support from the other vice presidents, Jon Bascom and Jim Etzkorn.

Q. Who is the last person you mentioned?

A. Jim Etzkorn.

Q. Could you please spell that name?

A. E-T-Z-K-O-R-N.

Q. Okay. And these were all vice presidents at AutoZone during that time period?

A. Around that time period. I really don't remember when either Jim Etzkorn started -- it was five or six or seven years ago -- or Bruce Clark, who also I believe started somewhere around the same time.

Q. Are you familiar with an individual named Jim Greer?

A. Yes.

Q. Did he play a role in this decision to migrate from OpenServer to Linux?

A. Yes.

Q. And were you his supervisor at that time?

A. No, I was not.

Q. Was he equivalent to you, or was he above you in the --

A. He was equivalent to me, same title, I believe.

Q. Now, you understood, did you not, that OpenServer was a proprietary operating system that AutoZone was licensing from SCO? Correct?

A. Yes.

Q. Do you know if anybody at AutoZone consulted an attorney, a copyright attorney, before deciding to migrate from OpenServer to Linux?

A. I don't know.

Q. To your knowledge, did they?

A. I don't know. I'm not aware of any consultation.

Q. Do you know if anybody reviewed the licenses with SCO to determine whether it was appropriate or

legal under those licenses to port binaries created for the OpenServer operating system to the Linux operating system?

MR. STEWART: Hang on just a second before you answer that. Korian, can I ask you to read that back?

(WHEREUPON, THE LAST QUESTION WAS READ BACK BY THE COURT REPORTER.)

MR. STEWART: I'm going to object to the question on the grounds that it's outside the deposition notice. But I'll allow you to answer. Answer just the question asked, though, so that we avoid getting off into any attorney/client issues.

A. I believe the answer is I don't know.

Q. (BY MR. STONE) Well, that makes it easy. Okay. So you don't have any knowledge about anybody doing that one way or the other. Is that fair to say?

A. I don't know whether we specifically investigated the porting of binaries to any operating system from any operating system.

Q. Well, then let me make that a more general question, because I didn't want to somehow limit it. Do you know if anybody at AutoZone examined the license agreements with SCO regarding OpenServer operating

system to determine whether the migration from OpenServer to Linux was appropriate or legal under those licenses?

MR. STEWART: Same objection as my prior objection. But you can answer.

A. I'm sorry. I still -- I don't know. I don't have firsthand knowledge of, or even secondhand knowledge, of whether or who might have investigated our SCO licensing with respect to the Linux port.

Q. (BY MR. STONE) Okay. Do you know what a developer kit is?

A. My definition of the developer kit is a set of software tools that we use to create the software or translate the software from source code to executable form. So it involves things like compilers, linkers, debuggers.

Q. Do you know whether AutoZone had licensed a developer kit from SCO when it was utilizing the OpenServer operating system?

MR. STEWART: Same objection. You can answer.

A. I believe we had developed a kit license.

Q. (BY MR. STONE) I'm going to ask you a question. Answer me in you own words. This is sort of what I'm getting at here. What I want to

understand is, AutoZone licensed the OpenServer operating system. That was an operating system that was on certain machines at AutoZone. Is that correct?

A. Yes.

Q. Then AutoZone created its own code, correct, that was going to then work on the OpenServer operating system? Is that right?

A. Say that again.

Q. I probably said it wrong. Let me try it a different way. AutoZone programmers created programs that would run on the OpenServer operating system. Correct?

A. Yes.

Q. You actually wrote your own programs at AutoZone. Correct?

A. Yes.

Q. And you wrote a lot of them. Is that right?

A. Yes.

Q. I mean, you have like 8 -- currently somewhere near like 8 million lines of code or something like that? Does that sound right?

A. I really don't know how much code we have.

Q. Okay. But it's -- it's a lot of them. Correct?

A. Right.

Q. Now, if I refer to those programs as binaries, is that also -- do you understand what I mean, or am I being technically incorrect by doing that?

A. The programs that we write are written in programming languages that are human readable.

Q. Which is source code?

A. Which is source code. The development tools turn that source code into binaries, which can then be loaded and run on the computer itself.

Q. So am I correct that you first write the program in some programming language? Correct?

A. Yes.

Q. And then that's what's called the source code?

A. Right.

Q. And then you create -- you either create or you're given some kind of compiler which is another type of software which turns that source code into a binary or an executable that the computer can execute? Is that right?

A. Right.

Q. Binaries are like zeros and ones? Is that right?

A. Right.

Q. Okay. So this is the part I need your help

with: When you go from a -- from the source code to a binary or an executable, what do you add -- do you add things to the source code that are not written by you? Do you add pieces of, for example, libraries or something else that allows the program to run on the particular operating system that you're dealing with?

A. You can link your -- the binary that's created from your source code with existing library code.

Q. Okay. So when you would write an executable or a binary for OpenServer, if it was the type of program that needed a library, you would then link it with one of the OpenServer libraries. Is that right?

A. That's correct.

Q. Well, now, would those libraries actually be in OpenServer, or would you have pieces of them that you could somehow with this developer kit link by using this compiler?

A. It actually can work either way. You can have library code that already exists on a running machine and then link with it at the time you try to run the program, or you can link code from the library directly into the executable that's created as a result of the compile.

Q. Is the first example that you gave me referred to as a dynamic library?

don't know if I could say that again.
(WHEREUPON, THE LAST QUESTION WAS READ BACK BY THE COURT REPORTER.)
A. Yes.
Q. (BY MR. STONE) Now, when you decided -- by you, I mean AutoZone, the group you've mentioned already at AutoZone -- to migrate to Linux, were you going to take these binaries that you'd written to work on the SCO OpenServer operating system and modify them so they would work on the Linux operating system?
A. No. We were going to recompile the source code to run natively on Linux.
Q. Okay. So you weren't going to take -- because the binaries, the final versions, ran on OpenServer. Correct?
A. Right.
Q. Am I right that it would be rare that you would be able to take a binary that ran on OpenServer and simply run it on Linux without doing something to it?
MR. STEWART: Object to form. You can answer.
A. I don't know. At the time, we believed that many SCO binaries would run natively on Linux. But we

weren't interested in, you know, risking odd behavior of running a non-native format on Linux.
Q. (BY MR. STONE) And when you're talking about a non-native format, are you talking about the COFF format versus some other format?
A. Versus whatever is produced when you develop something natively on Linux.
Q. For example, an ELF file? Would that be one?
A. That would be one.
Q. Okay. So you were going to take the source code that you used to compile these binaries that ran on OpenServer, and then you were going to compile it with a different compiler that would now put it in this format so it would run on Linux. Is that correct?
A. Right.
Q. But you were going to continue to use the programs that you'd written to perform essentially the same functions that they performed when you were using the OpenServer operating system. Is that right?
A. We were going to modify the source code to the extent that the Linux environment differed from the OpenServer environment.
Q. Okay. Let me try that question a different way. You'd written all these programs. There were like hundreds. Right?

A. (Witness nods affirmatively.)
Q. Were there thousands of programs?
A. Counting scripts and so forth, yes.
Q. And you weren't going to just rewrite all the programs and create a completely new way of doing business at AutoZone. Correct?
A. Right.
Q. The idea was to, with slight modifications, take those programs you had already written and use them on this new operating system. Correct?
MR. STEWART: Object to form. You can answer.
A. Yes.
Q. (BY MR. STONE) Okay. Now, you were involved also after this lawsuit was filed in reviewing what was actually on the AutoZone servers to determine whether SCO code or COFF files had been copied on to them. Is that right?
A. Yes.
Q. Were you the primary person responsible for doing that?
A. I was responsible for coordinating the identification of those files.
Q. Yeah. Let me clarify. I didn't mean that you went and did it all yourself. I meant, you were the top

person coordinating the people that were doing it. Is that correct?
A. Yes.
Q. Now, we've talked about a COFF file. Can you tell me what a COFF file is?
A. I believe it's a format for arranging the pieces of a program in its binary form.
Q. Okay. So -- and that would be a format that would work with a particular operating system that the complier was designing to work with. Is that right?
A. Yes.
Q. And was it -- am I correct that the COFF file format was a format that was used specifically for SCO licensed operating systems, to your knowledge?
A. I don't know about the -- I don't know how common a COFF file format is. I believe the acronym stands for common object file format, but I'm -- I'm not sure about that. So also, I don't know whether a COFF format -- COFF formatted file is proprietary format or not.
Q. Okay. Actually, I don't think I asked you whether it was a proprietary format. But let me try to rephrase my question because I think it may have been unclear.
AutoZone, when it wrote binaries in the COFF file

MR. STEWART: Object to form. You can answer.

Q. (BY MR. STONE) Would you refer to the first example that you gave me as a dynamic library or something else?

A. Typically, I would refer to it as a shared library.

Q. Okay. Is that because other programs can share the library?

A. Simultaneously.

Q. Okay. And the other example you gave me, would that be like a static library?

A. Yes.

Q. So when you use the word statically linked, that would be linking a static library to source code to create a binary to run on an operating system? Is that right?

A. Yes.

Q. And some of the binaries that AutoZone wrote to run on OpenServer were statically linked. Is that right?

A. Yes.

Q. So does that mean that the libraries that were statically linked were somehow separate from the OpenServer operating system, that they were on a, you

know, a CD or machine or something, and then you linked them so they became part of the program?

A. I'm not sure I understand the question.

Q. Probably because -- let me ask it so you can just answer it. When you link these static libraries, where do you get them from to link them to the source code that you were writing?

A. I believe the libraries are a part of the operating system distribution.

Q. So you could, for example, copy them into your binary or something like that? Is that what you're saying?

A. The linker --

MR. STEWART: Object to form. You can answer.

A. The linker would copy them into the object that it creates.

Q. (BY MR. STONE) So you'd run this -- this compiler. And as part of the compiler, it would find that library and copy it into your executable. Is that right?

A. On a typical compile with static library, that's right.

Q. Are you familiar with something called a header file?

A. Yes.

Q. Can you tell me what it is?

A. Header files are source code that are made available for developers' use in writing a software.

Q. So this would be source -- let me go back for a second. To your knowledge, did AutoZone have a source code license for all the source code to OpenServer?

A. No.

Q. So do I take your answer to mean that some of the sources -- you had a license to use some of the source code from OpenServer or that was proprietary to SCO, to use these header files in your -- creating your binaries?

A. Yes.

Q. And in some cases, you would use header files in the binaries when you did this compiling process? Is that right?

A. Yes.

Q. So that's another thing that you might link to your source code that you wrote to create a binary. I probably -- link is probably the wrong word.

A. Yes.

MR. STEWART: Yes, it's the wrong word or yes to the question?

Q. (BY MR. STONE) Yes, it's the wrong

word --

A. Yes to the question.

Q. Yes to the question. Okay. Other than libraries and header files, was there anything else that AutoZone licensed from SCO that you would have linked into the binaries when you create -- when you created them that wasn't written by AutoZone?

MR. STEWART: Object to form. But you can answer.

Q. (BY MR. STONE) If you don't understand it, I can --

A. I can't think of anything --

Q. Okay.

A. -- other than headers and libraries.

Q. Okay. So you understood that you had a license that allowed you to use these header files and these libraries to create your binaries to run on the OpenServer operating system. Correct?

MR. STEWART: I'm going to object to the extent it calls for a legal conclusion --

MR. STONE: Right.

MR. STEWART: -- but allow him to testify as to his knowledge or understanding.

A. I'm sorry. Say that again.

MR. STONE: Could you read it back? I

format, was writing them to operate on the open SCO licensed OpenServer operating system. Right?
A. Yes.
MR. STEWART: Object to form.
Q. (BY MR. STONE) Yes?
A. Yes.
Q. Okay. And different operating systems often use different file formats? Is that what you call them, file formats?
A. Yes.
Q. And Linux, you've told me, used a different file format, one of which was ELF. Is that right?
A. Yes.
Q. Do you know what ELF stands for?
A. I do not. Or if I know, I can't remember right now.
Q. Does L stand for linking?
A. I don't remember.
Q. Just stuck in my head. Okay. Now, would it be fair to say that if you didn't have the source code from which a binary was created, you couldn't recompile the binary in a different file format?
A. If you don't have the source code, yeah, you can't recompile for another environment.
Q. And some of the binaries that AutoZone used

were licensed from third parties, and AutoZone did not have access to the source code. Is that correct?
A. That's true.
Q. So it, therefore, couldn't recompile those binaries?
A. Right.
Q. Okay.
MR. STONE: Do you have the interrogatories?
Q. (BY MR. STONE) By the way, if you get tired and need to take a break for some reason or another, you know, you're not a prisoner here. You can let me -- let me know. Okay?
A. Thank you.
MR. STONE: Would you mark that as -- what's the next number?
MR. IANNICELLI: 33.
MR. STONE: Celmer 33. I'm just going in order of the exhibits, you know, since we marked 31 of them at Greer.
MR. STEWART: Okay. Great.
MR. STONE: Okay. But I'm putting the name on it so we know whose --
MR. STEWART: I was curious about that.
(WHEREUPON, THE ABOVE-MENTIONED

DOCUMENT WAS MARKED AS EXHIBIT NO. 33, AND IS HERETO ATTAChED.)
Q. (BY MR. STONE) Okay. I'll show you what I have marked as Celmer Exhibit 33, Mr. Celmer, and ask you if you've seen that before.
A. I think I have.
Q. I hope you have, because I think you're the one who gave the answers for some of these. But we'll get into it in a minute. I'll represent to you that these are responses to interrogatories and request for production that we served upon your counsel in this case. Take a look at Interrogatory Number 3.
MR. STEWART: It starts on top of Page 5.
Q. (BY MR. STONE) It starts on the top of Page 5. Correct. It says, States whether you or anyone on your behalf or at your direction copied all or any part of the SCO OpenServer operating system or any other Unix-based operating system or used any part of a SCO OpenServer operating system or any other Unix-based operating system with a Linux-based operating system, including but not limited to any SCO OpenServer for Unix-based software, source code, object code, libraries and/or programming tools, and identify any aspect of the SCO OpenServer operating system and any other

Unix-based operating system that was copied or was used with a Linux-based operating system by you or on your behalf or at your direction. Do you see that?
A. I do.
Q. Have you seen that question before?
A. I don't specifically remember seeing this question --
Q. Do you recall --
A. -- before.
Q. I'm sorry. Do you recall whether or not you were the person that was requested to look into this and answer this question?
A. In general, I'm not sure I was called upon to answer specifically any of these questions. Rather, I was given a set of tasks around gathering data that would provide the answers.
Q. Okay. Well, let's try it this way: If you take a look at the response, it says, AutoZone copied two third-party applications, Compx and Decompx, from its OpenServer installation on to its Linux installation. You see that?
A. Yes.
Q. Were you the person who went back and looked for Compx and Decompx on AutoZone's Linux installation?

A. I went and looked for any COFF files on our store system.
Q. And were Compx and Decompx two of the files that you found?
A. Yes.
Q. Okay. And then it goes on to say, because AutoZone does not have the source code for these programs, it cannot determine whether any OpenServer libraries were compiled with the software or not. AutoZone has not used Compx or Decompx programs since 2003. Do you see that?
A. Yes.
Q. Now, is that what we were talking about before, Compx and Decompx programs are third-party programs for which AutoZone didn't have access to the source code?
A. That's right.
Q. And these were programs that were written in the COFF format to work on OpenServer? Is that correct?
A. I believe they were COFF files, yes.
Q. And how did you determine that Compx or Decompx programs had not been used since 2003?
MR. STEWART: Object to form. You can answer.
A. I'm not sure I can place this date. But the

use of Compx and Decompx became unnecessary with an upgrade to our distribution center software.
Q. (BY MR. STONE) Do you know when that upgrade occurred?
A. I do not. I believe it's a 2003 date, but I don't know for sure when that happened.
Q. So you don't believe you were the person that checked specifically the date this upgrade occurred, as you sit here today?
A. I don't know if I provided that date. I do remember looking into what changes were made on the distribution center software side that made it unnecessary for us to use Compx and Decompx. And I certainly could have, you know, found this date and provided it. I don't specifically remember whether I did that.
Q. That's fine. Can you tell me what Compx and Decompx were intended to do, those programs?
A. These programs performed two primary functions. One is to compress and decompress data to make it easier to transfer and house. But the other feature they have is the translation of file formats, for example, from EBCDIC to ASCII or Big Indian to Little Indian, that sort of thing.
Q. So they can also help you translate formats?

Is that what you said?
A. Yes.
Q. And would it be fair to say you moved -- you moved -- AutoZone moved to Linux or migrated to Linux some time around 2000? Is that right?
A. I believe that's right.
Q. So from 2000 to 2003, Compx and Decompx were used by AutoZone on its servers. Correct?
A. Yes.
Q. And those were COFF files. Correct?
A. Yes.
Q. Those were files that you did not recompile in the Linux format. Correct?
A. That's right.
Q. Now, am I also correct that this statement here, you later learned to be inaccurate and in fact that Compx and Decompx were being used on at least some of your store servers after 2003?
A. Yes.
Q. And you learned that because when you tried to delete those files from your 3500 store servers, 650 of the machines had their replenishment system fail. Is that right?
A. I believe that's right.
Q. So I take it you concluded from that, that

Compx and Decompx were in some way being used by those 650 machines. Is that right?
A. That's right.
Q. Can you tell me when those files were deleted from your 3500 store systems?
A. I don't remember the dates.
MR. STONE: You know, why don't we just take a five-minute break.
VIDEO SPECIALIST: Off the record 9:53.
(SHORT BREAK)
VIDEO SPECIALIST: Back on the record at 10:05. This is the beginning of Tape Number 2.
Q. (BY MR. STONE) Mr. Greer [sic], would it be fair to say that when you were doing this migration to Linux for AutoZone and the people at AutoZone that were performing it, you weren't focused on the legal issue of whether you could use SCO libraries or code in the programs that you were porting to Linux? You were concerned with making the programs function efficiently on Linux.
A. That's true.
Q. And would it be also true that -- how many programs do you have working on this migration to Linux when you had -- at its height, let's say?
A. Well, it was a fairly long process, and we had

other things that we had to do. So I imagine at its height, there were 12 or 15 developers working on it simultaneously. But that -- that burst of energy, I don't know how long that would have lasted. It was more of a lengthy process with relatively few people cycling in and out of the project.

Q. Were these developers located in different locations around the country, or were they all in the same building or something else?

A. They all would have been in our building, I think.

Q. Was anybody monitoring to make sure that these various developers were not copying COFF binaries when they were doing the porting to Linux onto the Linux machine?

A. I don't know.

Q. You weren't doing it. Is that fair to say?

A. I was not doing it.

Q. I'm going to refer you back to an exhibit.

MR. STONE: I believe it's 32? Is that right?

MR. IANNICELLI: 33.

Q. (BY MR. STONE) 33, which is your -- the first set of interrogatory answers that we received from you in this case. And we already referred to

the statement about Compx and Decompx before the break. Do you recall that?

A. Yes.

Q. I'm on Page 5. You go on to say in this response, In preparing its response to this interrogatory, AutoZone discovered a few minor programs that were originally compiled on OpenServer had errantly been copied on to its Linux image, an image that contains approximately 700 programs. Is that right?

A. Yes.

Q. Did you later learn that it was more than a few minor programs that had been copied on to the store machines?

A. I'm not sure of the timing of this document. But we went through two investigative exercises, and we did find additional files the second time.

Q. Okay. I'm going to go through some of those with you and see if we can't understand what you did. Now, again, in this response to Interrogatory Number 3, you say that several of these programs are programs for the SCO store support group. Can you tell me what the SCO store support group is?

A. I believe this refers to our help desk which is the front line of technical support for our stores.

Q. So if somebody had a problem with the

operating system on the store server, they would call this support group and ask them their question?

A. Yes.

Q. Where was the support group located?

A. In our corporate offices.

Q. Is that in Memphis?

A. Yes.

Q. And would these programs be on the -- have been on the store servers, or would they have been on some other server or both?

A. Probably both.

Q. So these would have been programs that were utilized by SCO store support group in dealing with support questions from the field? Is that right?

A. Yes.

Q. And so if you copied these programs on to your store servers, it would have been on all 3500 machines? Is that right?

A. Yes.

Q. Then it says, And programs for the sorting module of the store management system. What is the sorting module of the store management system?

A. AutoZone uses a file access product that creates indexes into our data files. And in the terminology of that product, the creation of an index

file is called sorting because it makes this pile look as if it's been sorting.

Q. So this is something that creates indexes of the files that are on the store machines? Is that right?

A. That's right.

Q. And it says AutoZone has not used this since it converted to Linux. Was that an accurate statement or --

A. I believe that is an accurate statement.

Q. And what does AutoZone now use to create these indexes?

(BRIEF INTERRUPTION)

Q. (BY MR. STONE) Go ahead.

A. We still write individual programs that understand these data formats and create the indexes. I believe what this statement means is that these particular programs that create the indexes for these particular data files have not been used.

Q. Okay. So you still do -- you do something which performs this function, but you don't believe these particular programs have been used since converting to Linux?

A. I don't believe so.

Q. And how did you determine that they had not

been used since you converted to Linux?
A. Well, there are --
MR. STEWART: Object to form. You can answer.
A. There are two things. One is, these files -- let me see if I can remember correctly. There are really three issues. One is that these files very rarely change. And when I say very rarely, I mean, you know, they might change every several years.
Another issue is that when we change these files, we know they have to be changed everywhere because it's the same file in every store. And so what we typically do is include that file as part of what we call a release, just basically attached to our system. And so we send out this file as a -- we send it out intact. There's no need to run the program to create the index because we're sending the index to the store.
And then the third issue is I'm not certain these programs would run on our installation of Linux.
Q. (BY MR. STONE) Have you tried to run them?
A. I tried to run some of the programs that we found just out of curiosity. Many of them would not run. I don't remember specifically whether these would run or not.

Q. Did some of them run that you tried to run?
A. Some of them -- some of them tried to run. However, some of them were also pre-Y2K programs. So even had they run, they would not have been able to correctly interpret the data. I don't remember whether these were pre-Y2K or not.
Q. It says here, the next sentence, eight of the programs are still in sporadic use as part of the sorting module for AutoZone store management system. So there were eight programs that still were being used? Is that right?
A. The sporadic use probably refers to these same programs. Because in the event that we have file corruption, if it became necessary for the help desk to recreate the index, this is how they would have tried to do it.
Q. So would it be fair to say that you made a guess about whether or not, you know, an educated guess, about whether these programs have been used, but you have no way to know since you converted to Linux whether these programs were or were not used by servers at AutoZone?
MR. STEWART: And by you, do you mean the witness?
MR. STONE: I just mean the witness.

Yes.
A. If we try to run these programs and find out that they won't load or run, then clearly, they could not have been used. Barring that, it is -- if they will run and if they will run correctly, then it is possible that they may have been used.
Q. (BY MR. STONE) And you also made an educated guess that Compx and Decompx had not been used since 2003, and you turned out to be wrong about that. Correct?
A. That's correct.
Q. Okay. You go on to say that AutoZone has recompiled the programs in Linux, and AutoZone is now in the process of deleting from its store image all the OpenServer compiled programs that it no longer uses and replacing the eight OpenServer compiled sort programs with Linux compiled versions. And it says AutoZone anticipates that testing and rule out will be completed within one week.
Was it you who carried out these -- under your direction, carried out this work?
A. Yes.
Q. Okay. When you say AutoZone recompiled the programs in Linux, can you be specific about which programs you recompiled?

A. I don't remember the names of the programs, but they were the programs that were identified by our first investigation of COFF files.
Q. And how exactly did you go about recompiling them in Linux?
A. We simply checked the source code out of our source code repository on to our Linux development machine and invoked the GCC compiler to create the binaries for Linux.
Q. And did you have to link that source code with anything other than whatever that compiler linked it with?
MR. STEWART: Object to form. You can answer.
A. Yes. To the extent that we have libraries of code that understand the file formats themselves. So we had our source code which contains the logic for creating the index file.
Q. (BY MR. STONE) Right.
A. We have a repository of code that understands these data files formats, these index files and data files, that it was necessary to link with that code so that our logic would work. And it was necessary to link with operating system libraries so that the program would run at load time.

Q. And the operating system libraries that you linked with, were those ELF libraries?
A. Yes.
Q. Were any of them statically-linked libraries?
A. I don't think so. No.
Q. Do you know what a flag is?
A. A flag?
Q. Yes. I don't mean the one that you --
A. Right.
Q. -- raise up and salute.
A. In common usage, it's a small indicator of something. It typically is an on/off kind of deal. Some functionality is enabled or disabled.
Q. When you write a source code to create a binary, do you put flags in it?
A. Yes.
Q. And what is -- what is the purpose of those flags?
A. There can be many purposes for flag. You can allow the program itself to take some action or not, depending on the condition of that flag, which represents some business meaning. You can also create the ability for the compiler itself to determine which code should be compiled and which code should not be compiled into an object.

Q. So you could put in a flag that tells the compiler something. Is that right?
A. That's right.
Q. To do or not do something?
A. That's right.
Q. When you recompile these programs to run on Linux, do you have to go in and change flags in the source code because you were using a different compiler?
A. Sometimes the -- either the program logic or the compiler itself or the operating system will demand different things. And the presence of these flags allows us to have one source code that is understandable by either compiler.
Q. I'm trying to figure out what you just said. You said the presence of these flags allows the -- allows the operating system? Say it again. I'm sorry.
A. Just as the source code describes a set of actions that the program should take --
Q. Right.
A. -- you can use these flags to instruct the compiler which pieces of source code it should compile versus which pieces of source code it should ignore.
Q. Okay. So for example -- I think I understand this now. So you could write source code which might be compiled to work on Linux or work on OpenServer but just

have different flags to tell the compiler which pieces of the source code to use? Is that right?
A. That's right.
MR. STONE: Can you mark this as the next exhibit?
(WHEREUPON, THE ABOVE-MENTIONED DOCUMENT WAS MARKED AS EXHIBIT NO. 34, AND IS HERETO ATTACHED.)
Q. (BY MR. STONE) Can you take a look at what's been marked as Exhibit 34. It's an e-mail that was produced to us by your counsel, which we believe is e-mail relating to the porting process. But all I really want you to do is read it. And if you can explain to me, explain to me what these people are talking about in this e-mail.
A. I think this e-mail is a suggestion by one of our contractors for a way to simplify the port.
Q. When you say contractors, what do you mean contractors?
A. Just one minute.
Q. Sure. Take your time.
(Brief pause)
A. Okay. I'm sorry. I remember either this e-mail or one very much like it. Because there was some concern about how we were going to manage to compile

things two ways. The context for this one, though, is the Y2K project. We needed to be able to compile our code in pre-Y2K format in order to support our existing operation and also in post-Y2K format in order for the preparation of the rollout of the software and data that would be used subsequent to Y2K.
Q. And that's what you believe this e-mail is talking about, how to accomplish that?
A. That's right. We wanted to make things as simple for the developers as possible. Because it's much easier to have the system always work the same way rather than educate our entire staff to do things one way under a certain set of circumstances and a different way under another set of circumstances.
And so we were trying to figure out how we could allow our developers to do nothing more than issue the same make command that they always have issued in the past but allow the system to intelligently determine what kind of compile they really needed to do and which set of AutoZone include files they needed to use for that compile.
Q. Let me ask you a question about what you said in your -- I think it was the previous answer to this answer. Am I correct that with some of these programs that you recompiled in Linux, rather than rewriting the

source code, you added source code and changed flags or added flags so that the source code could run both on OpenServer -- I mean, they could be compiled on both OpenServer and Linux? Is that right?

A. That's right.

Q. So you would go in and you'd make these modifications to your source code, and then the compiler would -- it would tell the compiler take like this part if you're going to do it on Linux and take this part if you're going to do it on OpenServer. Is that right?

A. That's right.

Q. Did you do that with most of the programs, or was that just only in a few rare cases or something else?

A. I think that probably half of our programs, just programs, needed those kinds of flags. That's just a guess, though, really.

Q. And you don't remember how many programs it was that you actually ported?

A. I'm sorry. I don't remember.

Q. Do you remember like less than a thousand, more than a thousand?

A. I'm sure it was more than a thousand if you include scripts, which I do because we had to port them.

Q. Okay. When you include scripts, was it more

than 2,000?

A. I don't think so.

Q. So your --

A. I really don't know.

Q. -- best recollection would be somewhere between a thousand and 2,000 if you include scripts. Is that right?

A. Yes. But I'm not very confident in that answer.

Q. I understand you're giving me the best of your recollection. Could you move to Page 7 of Exhibit 33, please. You see there it says interrogatory number -- by the way, if you see Interrogatory Number 6, it says identify each person who provided the information for use in answering these interrogatories. And the response is Bob Celmer. So if you hadn't seen these before, I would be a little concerned. But you believe you either saw them or the substance of what was in them was communicated to you. Is that fair to say?

MR. STEWART: I object to the question because that's not a complete response to Interrogatory Number 6. But subject to that objection, you can answer.

Q. (BY MR. STONE) That's true. It also says -- I want to be fair. It says, additionally,

certain information used to respond to these interrogatories is derived from information provided to SCO by Jim Greer during his deposition.

A. I'm not sure whether this is the specific document I have seen before. But these questions and responses are not unfamiliar to me.

Q. Okay. So Interrogatory Number 7 says, state whether you have used any static-shared libraries and any dynamic-shared libraries, and if so, identify each of them. Then it goes on to ask other questions about them. Do you see that?

A. Yes.

Q. Then there's a bunch of objections. And it says, Subject to and without waiving the foregoing objections, AutoZone states that it has not used OpenServer static-shared libraries or dynamic-shared libraries based upon definitions of those terms contained on SCO's website. Do you see that?

A. Yes, I do.

Q. Do you now know that statement to be inaccurate?

A. I now know that Compx and Decompx were being used when I previously did not believe that they were being used.

Q. Okay.

A. And if Compx and Decompx contained library code, then that exception -- yeah, that statement would not be true.

Q. Wasn't AutoZone using some COFF sort files that were loaded on to its Linux image for some period of time?

A. Are you referring to the eight that we talked about a minute ago?

Q. I don't know because it's confusing because I got three different responses from you and AutoZone, and it's hard to tell whether they're referring to the same responses or not. So I can go through the responses, and we'll try to clarify that. But I guess my question stands. And you know, whichever ones you're referring to, just tell me, and you can answer it that way.

A. I'm sorry. Would you repeat the question, please.

Q. That there were nine COFF sort files that were loaded on to the Linux image that were on 3500 of your machines until they removed -- they were removed some time in October of 2004, that were being used in some form.

A. I think that refers to the same files that we were just talking about earlier. And I don't know that that or whether those files were being used.

Q. You just don't know one way or the other whether they were being used. Is that correct?
A. I don't know one way or the other whether they were being used. But I do know that they were on the image machine and that's how it got propagated to all those stores.
Q. And do you also know whether any of those files contains statically-linked libraries?
A. I believe those programs were statically linked under OpenServer.
Q. So if they were being used, then AutoZone would have been using statically-linked libraries. Correct?
A. Yes.
Q. Okay. Do you not -- I'm trying to understand the statement in here about static-shared libraries. Do you not use the term static-shared libraries?
A. I have not used the term static-shared libraries.
Q. So what do you call those libraries that you -- that AutoZone linked statically to its source code when it was creating OpenServer binaries?
A. I just call them static libraries.
Q. Okay. So when you were answering this question, were you answering it based on some

distinction you were making between a static-shared library and a static library?
MR. STEWART: Object to form. You can answer.
Q. (BY MR. STONE) Or were you talking about what you understood to be the libraries that AutoZone had linked when it created binaries to run on OpenServer?
MR. STEWART: Same objection.
A. I frankly don't know because I've had to change my definition of what this term means. I don't know when that change came about, to be honest with you.
Q. (BY MR. STONE) Well, would it be fair to say that when you first responded to discovery request and did what you call your first run through of the AutoZone files to determine what COFF files might be there, you believed that no static libraries that you had linked when you created these binaries to work on OpenServer had been used in the Linux operating system at AutoZone.
A. That's right.
Q. And you later found out that that may not be true as to Decompx and Compx, and you don't know whether it's true as to the nine COFF sort files that we've been discussing. Is that correct?

A. That's correct.
Q. Okay.
MR. STONE: Mark this as the next exhibit, please.
(WHEREUPON, THE ABOVE-MENTIONED DOCUMENT WAS MARKED AS EXHIBIT NO. 35, AND IS HERETO ATTACHED.)
Q. (BY MR. STONE) Mr. Celmer, I'm showing you what I've marked as Exhibit 35, which is denominated as AutoZone's responses to our first set of requests for admissions. Actually, to Defendant AutoZone, Inc. and supplemental interrogatory. So I guess I guess it's both. You were answering our requests for admissions, and you were also responding to a supplementary interrogatory that we had sent you. Do you recall seeing this document before?
A. Yes, or one just like it.
Q. Okay. Refer you to Page 5. Request for admission number 10 requests, Admit that you copied, referenced, or relied upon some portion or part of the Unix-based operating system as part of your conversion from the use of a Unix-based operating system to the use of a Linux-based operating system. You see that?
A. Yes.

Q. And then the response was, AutoZone denies that it referenced or relied upon any portion of a Unix-based operating system as part of AutoZone's conversion of its store servers to Linux. AutoZone admits that it errantly copied a few minor programs to its Linux store image that were compiled under OpenServer and therefore included some code from five basic OpenServer static libraries. AutoZone is in the process of deleting or recompiling all of these files. Process should be completed within one week. Do you see that?
A. Yes.
Q. Would it be fair to say that you later learned that there were other files that had been copied on to AutoZone's Linux operating system that you did not know about when you responded to this request for admission number 10?
MR. STEWART: Objection. Asked and answered.
Q. (BY MR. STONE) You can answer.
A. Yes.
Q. And in fact, you found that there were 1681 COFF files on 387 AutoZone store machines at a later date. Correct?
A. That sounds about right.

Q. Now, this reference to the few minor programs, are we again talking about the 19 unused COFF files and 9 used COFF files that you advised SCO about, or are we just talking about the 8 or 9 files we've been talking about previously?

MR. STEWART: Object to form. You can answer.

A. I think this refers to all of those.

Q. (BY MR. STONE) So did you go through and determine that some of these files had OpenServer static libraries statically linked to them?

A. I think so.

Q. I'm just referring to the statement which says it included some code from five basic OpenServer static libraries. And I'm asking you were you the person who either determined or caused someone to determine that that was the case.

A. I did investigate what might have been linked into these programs. I'm not certain that this response is a result of those investigations.

Q. Because you don't recall what point in time that happened with respect to when this response was sent to us.

A. That's right.

Q. Okay.

MR. STONE: Just for the record, this response is dated October 20th, 2004.

Q. (BY MR. STONE) Can you tell me what the five basic OpenServer static libraries that are referred to in request for admission response Number 10 are?

MR. STEWART: Object to form. You can answer.

A. No. I'm sorry. I don't know what they are.

Q. (BY MR. STONE) Is that written down somewhere to your knowledge that -- did you go through and write down the different OpenServer static libraries that you determined were linked to files which had been copied on to the Linux AutoZone server?

MR. STEWART: Object to form. You can answer.

A. I don't know. But because we've quoted an exact number here, it's likely that it was written down some place.

MR. STONE: I'll make a request, and we'll follow it up with a letter to get that information of which libraries are being referenced.

MR. STEWART: And just for the record, my silence is not an indication that we will provide

those. We'll review it when we get your letter.

MR. STONE: As we've done in the past, I'll make requests; I'll follow them up with letters. And I presume that your silence does not mean that you're necessarily going to give me what I'm asking for.

MR. STEWART: Fair enough.

MR. STONE: Okay.

Q. (BY MR. STONE) When you -- when you went through and found out that these programs had been copied on to the Linux image, could you determine who actually copied them?

A. Sometimes, yes; most of the time, no.

Q. How could you determine?

A. I remember there were some small number of programs, you know, two, three, four, something like that, that were copied into the Linux image by mistake as part of the creation of one of our software releases. The -- I don't remember the name of the application. But at that time, we were having to create two releases, one for our stores and one for a small set of Linux test stores.

And whatever this release was called, there was another version called Linux, whatever that name is. And it was supposed to have the Linux versions of those

programs, and it had the SCO version of those programs. And so that release was sent to the stores from which we made -- subsequently made the images.

Q. And that's the 3500 stores?

A. As I remember it, it was during our test period where there were only four stores.

Q. So you believe that release was sent to the four stores?

A. That's right.

Q. And how many -- how many COFF files were in that release?

A. I believe there was a number like two or three or four.

Q. If you'll turn to Page 7. You see it's a -- request is that you admit you copied binary files from a SCO OpenServer operating system. And your response was, AutoZone denies copying binary files from a SCO OpenServer operating system to a computer running in Linux-based operating system except inasmuch as binary files from an SCO OpenServer operating system were statically linked inside errantly copied non-operating system binary files with eight minor sort files. Do you see that?

A. Yes, I do.

Q. And this is again the response that we also

referred to in your answer to interrogatories that you later learned to be inaccurate, correct, that there were additional files that you did not include in this response. Correct?

A. There were additional files we didn't know about. That's right.

Q. And then if you look at request for admission Number 21, which is on Page 8 at the bottom, and your response on Page 9, am I correct that again, we're referring to those eight files that we've discussed at some length in this response and at that time, you didn't realize that there were other files that you later located. Is that right?

A. That's right.

Q. So this admission -- this response to this admission is inaccurate. And today, as you sit here, you would modify it. Correct? It's incomplete is a better way of putting it.

A. It is incomplete.

Q. Can you explain to me what computers exist at AutoZone, where they're located, and what operating system they have on them? I mean, is everything at -- and let me be clear what I mean by that. Is everything at AutoZone on a Linux operating system, or are there computers that use other operating systems, and if so,

what are they?

A. We have a fairly broad IT organization. So we have quite a lot of computers. For example, we have desktop computers, which might be diskless workstations. They may have an imbedded operating system. We have Windows-based PCs on desktops. We have Windows servers. We have IBM servers running AIX. We have a mainframe computer running IBM's MVS operating system. We have Linux servers. We have a number of computers that serve security functions such as firewalls. And of course, we have the stores which are running Linux.

Q. So you've got the 3500 stores that are out there, and they have a serve -- each store has a server in it which is running Linux? Is that right?

A. That's right.

Q. And do they then have workstations that work off of that server, or is it just that one server that's in the store that really runs the Linux?

A. The hardware layout of our stores is changing. Historically, we had a server which ran all the programs. And attached to that server were dumb terminals and a few cash registers which themselves were some variety of PC with extra circuitry to support cash drawers and display bars and so forth. Those machines ran DOS. And all they did was operate the hardware and

run a terminal -- terminal immolator. The program itself was run on the central server.

That arrangement is changing today in that we do have more thin clients and fewer dumb terminals and -- struggling for the name of the device -- hand-held computers that use wireless networking and connect with the server.

Q. You mean, trios or something or --

A. Actually, I believe Symbol Technologies makes the device. But they are a PDA like device.

Q. So you use some kind of satellite to send programs out or to send images out to these Linux stores? Is that right?

A. The network connectivity that we have with almost all of the stores is satellite.

Q. And how long have you had that?

A. It's been a long time. Almost ten years, probably.

Q. So you had it when you had OpenServer. Correct? When you were using OpenServer as your store -- on your store machines?

A. Yes.

Q. And so when you would send images to your store machines or new releases, you would do it with this satellite? Is that right?

A. That's right.

Q. So you would create it in your -- your IT department would create it in Memphis, and then you'd beam it out over the satellite to the various stores?

A. Right.

Q. Did anybody ever go to the stores to then see whether the servers were receiving it correctly, or how did you determine that it worked?

A. Our release process involves a number of steps. So when you say how did you verify whether it works, it's sort of a long answer because it's a complicated process.

We begin with the developer who assembles all of the components of this -- that release and then puts them on to one of our test stations. Our test stations are built very much like a store, so that we're able to test what the developer believes to be his entire package.

The release -- skipping forward to just the use of the release system in the satellite, what we would typically do is transmit that package by satellite to a store where part of that package is a script that does the installation, the installation being the copying of the executable programs into their production locations along with whatever other steps were necessary like the

creation of data files or addition or deletion of data so that the system itself would work once the script completes.

So as long as we have a script, we're able to allow the script to be intelligent and tell us about errors. But it is incumbent upon the developer typically to manually verify the proper operation of the software or at least review error logs that are produced in the stores to see whether errors occurred.

Q. Am I correct that if a program is copied on to the Linux stored image that is in your, you know, IT department and beamed out over this satellite to the stores, that essentially 3500 copies of that program will then be copied on to each one of these servers that's in each one of the 3500 stores, in other words, one on each of the 3500?

A. While it's possible to use the satellite technology in that way, we typically don't as a way of risk management. We would not want to make a mistake in every store at the same time.

Q. So how do you do it, then?

A. Typically, we begin, as I mentioned before, with a single store or a very small group of stores. And then we expand the test, and we further expand the test. And our typical rollout would involve one week of

increased beta testing up to, say, a hundred stores or 200 stores. And then we would take an additional two weeks to make that change in all of the rest of the 3500 stores.

Q. So you do it -- you phase it in?

A. (Witness nods affirmatively.)

Q. But at the end of this process, these -- this file, let's say, if it was in the store image and you decided that it was working and you were going to use it, there would be one copy of that file in each of these 3500 servers. Is that right?

A. If that change was to be sent to the entire chain, that's true.

Q. So would that be true for, for example, Compx and Decompx that we talked about before?

MR. STEWART: Object to form. You can answer.

A. That would be true for Compx and Decompx if we were going to send Compx and Decompx out to our stores. Then it would follow the release pattern I've just described.

Q. (BY MR. STONE) Well, didn't you send Compx and Decompx out to your stores?

A. I don't know for sure. But I believe that Compx and Decompx were installed on the initial four

stores, and that subsequently -- subsequently, we created an image from one of those stores, and that image was eventually broadcast to every store.

Q. So when you went back pursuant to this process that you've described in your discovery responses to us, you found that Compx and Decompx were copied on every one of the 3500 servers that you had in your stores, and that's when you deleted them. Is that right?

MR. STEWART: Object to form. You can answer.

A. That's right.

Q. (BY MR. STONE) And that's when you deleted them and you found that 650 of those machines had their replenishment system fail. Correct?

A. That's right.

(WHEREUPON, A DOCUMENT WAS MARKED AS EXHIBIT NO. 36, AND IS HERETO ATTACHED.)

Q. (BY MR. STONE) I'll show you what's been marked as Exhibit 36 and ask you if you've seen that document before. It's a letter dated November 24th, 2004 from Mr. Stewart to me.

A. I believe I have seen it. Yes.

Q. Okay. Before I ask you about the letter, just

getting back to what you were telling me about the Linux systems and where they were, there's one in each of the stores. Correct?

A. That's right.

Q. Then are there -- how many Linux servers are there in your headquarters?

A. A dozen, perhaps.

Q. And --

A. No. There would be a dozen plus some number of test machines that also run Linux, which may be, I don't know, another half dozen, probably somewhere in the list of ten.

Q. So somewhere between 12 and 24 if you include the test machines? Is that right.

A. Sounds reasonable. Yes.

Q. Did you look at those machines when you were looking to see if COFF files had been copied on to Linux operating systems?

A. Yes.

Q. So you looked at those however many it was machines, Linux operating system machines that were in the headquarters, as well?

A. Yes.

Q. And would that include -- that included Spirit and Wrangler? Is that right?

A. Yes.
Q. Other than Spirit and Wrangler, did you find COFF files on any of the other machines in your headquarters?
A. Other than Spirit or Wrangler. I really don't remember specifically where we found COFF files. Seems that we may have -- I just don't know. We may have found them on other machines. I just don't remember right now.
Q. If you found them on other machines, did you produce discovery to us which would allow us to determine what the files were and what machines you found them on?
A. Yes.
Q. Am I right that Spirit and Wrangler were machines that you used when you -- when AutoZone made the migration to Linux?
A. Yes.
Q. And am I correct that what you did is you -- you had the Spirit machine was running OpenServer and the Wrangler machine was running Linux?
A. Right.
Q. And that Mr. Greer told us -- and that's why I'm saying it -- you can tell me if you understand that to be the case -- that you were running these two

machines simultaneously and using them to help port the source code that had been written for the OpenServer system to the Linux system. Is that right?
A. Wrangler was the development machine on which we did compiles and editing of source. Spirit was the source code repository. So we would have checked code out of the repository from Spirit on to Wrangler, made our changes on Wrangler, and then put those changed files back on to Spirit to safe keep source.
Q. And Spirit also had SCO licensed code on it, as well. Correct? It didn't just have your source code?
A. It was an OpenServer machine. Yes.
Q. Right. And you kept those machines till -- to the date of filing of this lawsuit. And when you looked as those machines, you found that there were COFF files that that been compiled for OpenServer on both those machines. Correct?
A. That's right.
Q. And I'll get into it in the letter. But you found that. Correct?
A. Yes.
Q. Is Spirit currently being used for any purpose by AutoZone or was it prior to the filing of this lawsuit?

MR. STEWART: I'm sorry. Which question?
MR. STONE: Well, let's start with the first one.
Q. (BY MR. STONE) Is it currently being used for any purpose?
A. Yes, it is.
Q. What is it used for?
A. It is still our source code repository.
Q. So this is the source code that you wrote during the time that you were licensing the OpenServer software that you used to compile binaries to work on the OpenServer software stored on the Spirit server. Is that right?
A. Yes.
Q. And then there's additional source code you've written since that time since you migrated to Linux? Is that right?
A. Yes.
Q. Is that -- yes?
A. Yes.
Q. Is that also stored on the Spirit server?
A. Yes.
Q. And is the Wrangler machine currently being used?
A. Yes, it is.

Q. What is it being used for?
A. It is still a Linux development machine.
Q. So was it used as a Linux development machine from the time that you first started using it when you began the porting process until this date, essentially?
A. Yes.
Q. Have you heard of a machine called Vision?
A. Yes.
Q. Tell me what that is.
A. It was the development machine for SCO development.
Q. So that was the equivalent of Wrangler for Linux? Is that right?
A. That's right. It was the machine that developers used while we were using SCO.
Q. And does Vision still exist?
A. Yes, it does.
Q. And is it still used for any purpose?
A. It doesn't get much use. But the files that I made copies of before deleting them from our stores are housed on Vision. And it probably gets occasional use but not much.
Q. Do you know if Vision was used during the porting process by AutoZone?
A. It was definitely in use during the porting

process because we were still using SCO in our stores. And so we were still doing development on that machine.

Q. Did you look at the Vision machine -- I guess you're telling me the Vision machine had the OpenServer operating system on it, then?

A. That's right.

Q. So it had SCO license code on it, as well?

A. Yes.

Q. And have you produced to us whatever code you found on that Vision machine? Do you know?

A. I doubt it.

Q. Okay. So the code that you produced to us is the only code that you found on Linux server machines and on Spirit. Is that right?

A. That's right.

Q. All right. Let's take a look at Exhibit 36. You recall that after the previous exhibits we looked at, you went back and you looked further at what OpenServer compiled programs might be in the Linux operating system at AutoZone?

A. Say that again.

Q. We've discussed, you know, the request for admissions and the interrogatories. Correct?

A. Right.

Q. And we went through the answers that you gave

on behalf of AutoZone in those. Correct?

A. Right.

Q. And you agree that you later learned that those were incomplete. Correct?

A. Right.

Q. And in some cases, they were inaccurate where you, for example, denied using a program which you later learned was used, for example, Decompx and Compx. Correct?

A. Right.

Q. What I'm talking about now is that this Exhibit 36 was the document by which you were informing us of this additional information that you'd learned since the previous exhibits that we've discussed. Okay?

A. Okay.

Q. And what I was asking you is, in order to come up with the information in Exhibit 36, did you go back, and did you look at the store servers and the servers in AutoZone to discover whether there was additional COFF files that were compiled for OpenServer on Linux machines?

A. Again, I'm not sure of the timing. But we have -- after the initial release that identified and removed the COFF files on the very first round, we did go back to the stores just to verify that they were all

gone, and we went back through the Linux servers in the corporate offices, also.

Q. Okay. All right. Well let's just look at the document and go from there. It says in the second paragraph under the word store servers. Do you see that?

A. Yes.

Q. It says, On October 19th, 20th, and 21, AutoZone sent the release of its software to its store servers that updated the store servers to replace 9 COFF files with Linux files and to delete 19 unused COFF files. Do you see that?

A. Yes.

Q. So does this refresh your recollection that there were 9 COFF files that were used up until that date in the store servers and that there were 19 files that you believed were unused?

MR. STEWART: Object to form. You can answer.

A. There were 9 COFF files, which -- let me start at the other side. There were 19 that we were never going to use. They were just useless files. There were 9 that we might some day need. Therefore, they were not obsolete programs. So we needed to replace them in case they were needed.

Q. (BY MR. STONE) Okay. Let me -- with that clarification, let me ask the question again. When you referred to the 19 unused COFF files, you were confident that they hadn't been used based on your review of those files. Is that correct?

A. That's right.

Q. With respect to the 9 COFF files referred to here, we discussed this before, you're not sure, as you sit here today, whether or not they may have been used at some time. Is that correct?

A. That's right.

Q. Okay. And you had to rewrite them because they were needed for -- well, you had to do something with them -- I don't want to use the wrong term. But you had to do something with them because this type of program was needed for your stores. Is that right?

MR. STEWART: Object to form.

A. It might be needed for the stores. In particular, I keep saying might because of those sort programs.

Q. (BY MR. STONE) Okay. And these nine programs, you recompiled for Linux. Is that right?

A. That's right.

Q. And are these programs that as we discussed, you went in and modified the source code so it would

work both on -- would compile both for Linux and OpenServer, if you recall?

A. I don't recall whether any of them had to be modified, but then, I didn't change them myself. It is possible that we had to put in some flags in order to allow Linux and SCO compiles. I just don't know for sure.

Q. And the 19 files, you deleted using this satellite? Is that right?

A. That's right. We used the satellite to distribute the release which contained the instructions to find and remove the files.

Q. Okay. So you're here today to tell me that the nine COFF files that you located are no longer being used by AutoZone. Correct?

MR. IANNICELLI:

MR. STEWART: Object to form. He's here today to testify in response a 30(b)(6)deposition notice.

MR. STONE: He can answer the question.

MR. STEWART: You can answer the question.

A. We replaced the nine files because they might be needed.

Q. (BY MR. STONE) Right.

A. The others were deleted.

Q. The 19 were deleted, and the other 9 were recompiled with a Linux compiler. Is that right?

A. That's right.

Q. And can you represent to me, as you sit here today, that none of those files contain any licensed SCO code or libraries?

A. I don't believe any of those files contain any licensed SCO.

Q. Okay. Then you go on in the letter to say, Two other COFF programs, Compx and Decompx, were also deleted in this release. You see that? We've talked about Compx and Decompx.

A. Yes.

Q. So this was when you deleted them by using the satellite. Is that right?

A. That's right.

Q. And you deleted them from all 3500 stores.

A. That's right.

Q. Okay. Then you say the nine replaced COFF files were on the store servers because they were inadvertently included in a Linux only release in November of 2001. Is that what you had referred to earlier in this deposition?

A. I believe so.

Q. Okay. So I think you earlier told me that there were -- there was a Linux release, and there was some other kind of release, and somehow, the COFF files got in the Linux release. Is that right?

A. Right.

Q. As you sit here today, do you recall how many machines those nine COFF files were on before they were deleted?

A. I believe they were on all of our store servers.

Q. Okay. So all 3500?

A. Right.

Q. Okay. You can see now why -- I don't mean this in a derogatory. But what's confusing about this is that you go on and you talk about these 19 unused COFF files. You say most were help utilities used outside of AutoZone store management system that were errantly omitted from the original conversion schedule and loaded into the Linux load image. Do you see that?

A. Yes.

Q. Is this again referring to the file -- some of the files that you have already told us about, the eight or nine files that were in your earlier response? That's what I'm trying to understand. Or is this 19 new files?

A. I think that throughout these documents, we're referring to two sets of files. Those were -- that were deleted because we simply didn't need them and the smaller set that we chose to replace because we might need them. So these 19 consisted of a number of different kinds of programs, some of which were used by our help desk when we were using SCO. So they were SCO related utilities.

Q. So are these 19 COFF files, files that you discovered after you answered our request for admissions and after you answered our answers to interrogatories that are in addition to the 9 files that you discussed in your answers to admissions?

MR. STEWART: Object to form. If you need to go back and take a look at the exhibits, go ahead and do so.

MR. STONE: Yeah, sure.

Q. (BY MR. STONE) This is where my confusion arises. So if you could clarify it for me, I'd appreciate it. Let me see if I can help you. If you look at admission Number 17, Page 7.

A. Okay.

Q. Do you see it says SCO OpenServer operating system worst -- I'm sorry. Binary files from SCO OpenServer operating system were statically linked

inside errantly copied non-operating system binary files or eight minor sort files, as described in response to request Number 10 above. So it talks about eight files.

And actually, request Number 10 doesn't talk about eight files. So this is the only place that I saw eight files.

So is this eight files the nine files? Is this eight files some other files? I'm trying to understand how I organize these files.

A. I believe that the eight files in this document are intended to refer to the same files that the nine files in this document are intended to refer to.

Q. Okay. So the eight files became nine files, and then you found an additional 19 files. Is that right?

A. I think we found the eight or nine and the 19 all at the same time.

Q. Okay.

A. The distinction between them is not in the timing but in what we did about them.

Q. Oh, okay. So you didn't believe -- you didn't include those 19 in these responses because these responses were requesting information about using static-shared library files, and you'd concluded that

you hadn't used those 19? Is that --

MR. STEWART: Object to form. You can answer.

A. That doesn't sound right.

Q. (BY MR. STONE) We never learned of these -- to my knowledge, we didn't learn of these 19 files until we got this November 24th letter.

A. Okay. Let me think for a minute.

Q. Okay.

A. I think I remember what happened.

Q. All right.

A. I think what happened was that I originally identified, by checking a small number of stores, a set of COFF files. How confident am I in this answer?

Q. It's okay.

A. I'm having trouble --

Q. Yeah. It's okay.

A. -- remembering the order.

Q. So am I correct that as you sit here today, you don't remember exactly which eight files were referred to in this response to request to admission but that I can be confident that the 19 and the 9 referred to in the November 24th letter include those 8 files?

A. Yes.

Q. Okay. And you believe it's within the nine

files. Is that right?

A. I believe so.

Q. Okay. And then if you look -- you turn the page of Exhibit 36, it says it is likely that many of these programs would not run under Linux due to compatibility issues for the fact that these programs predated year 2000 and would not have operated correctly with the new Autozone data format in use after the year 2000. Do you see that?

A. Yes.

Q. That's actually why I asked you at the beginning of the deposition about whether COFF files would run on Linux. Is this -- in your understanding, is this an accurate statement that it's lucky that -- it's likely that many of these COFF programs would not run under Linux due to compatibility issues?

MR. STEWART: Object to the form. That's not what the statement in this letter says.

MR. STONE: Well, I read the whole statement.

MR. STEWART: Well, the question you just asked, though, is only part of that statement.

MR. STONE: Right. I understand that.

Q. (BY MR. STONE) There's also the issue of this 2000 format. Correct?

A. Right.

Q. That you -- that you brought in for the Y2K.

A. Right.

Q. I guess what my question to you is, is there something about these binaries that would make them -- putting aside the Y2K issue that you told me about, you said that probably would make it difficult for them to run. Is there something about these particular binaries that you looked at them and determined that they wouldn't run under Linux due to compatibility issues, or was that a statement about OpenServer binaries in general?

MR. STEWART: Object to form. You can answer.

A. It was a statement about these programs in particular.

Q. (BY MR. STONE) Okay. What was it about them that you saw that suggested to you that they would not -- they would have compatibility issues?

A. I believe that I tried to run some of these programs just to see what they were or what they did. And some of them exited with an arrow that they couldn't be loaded. Some of them immediately died with an error. As opposed to printing an error message, they just died with an error.

And then others, I already knew what they did, and the files that they were intended to use either were obsolete and no longer exist on our servers or the dates on the files were for Y2K conversions. So they could not have used our updated data files structures.

Q. Okay. Then you go on here and you say in addition to the help utilities, there were four SCO compiled programs that were copied to the Linux machines errantly because the programs were not located in the proper directory for binary files prior to Linux conversion. Correct?

MR. STEWART: Objection. And just for clarification, Mr. Celmer didn't write this letter. I did.

MR. STONE: I'm sorry. I meant AutoZone goes on to say.

Q. (BY MR. STONE) But you said you've read the letter and you were part of the process of compiling information. Correct?

A. That's right.

Q. Were you the one that was responsible for determining that these programs were not located in the proper directory?

A. Yes.

Q. Just because they weren't in the proper

directory wouldn't mean they wouldn't run. Correct?

A. I don't know. The reason I don't know is that if you were to invoke them from the command line, they might run. However, our stores don't have access to a command line. They run software from the menu. And just by invoking the program name, you wouldn't get these. You would get one of the files that was in its proper location.

Q. Now, are these part of the -- this is all a description with regard to the 19 unused COFF files? Is that right? I'm just trying to understand what this is a subset of that you're describing here.

MR. STEWART: Again, same objection.

Q. (BY MR. STONE) That AutoZone is describing in its letter.

(BRIEF PAUSE)

A. Okay. I just reread this paragraph. And I'm sorry. But can you ask the question?

Q. (BY MR. STONE) Sure. Actually, you might want to look at the first two sentences of the previous page that says, of the 19 unused COFF files, and then it goes on to say most were help utilities used outside of AutoZone's store management system. And then you go on with the rest of this paragraph. So it was my understanding,

unless you're going to tell me differently, that this refers to the 19 unused files. Is that right?

A. I believe that's right.

Q. Okay. And the next paragraph goes on to say that deletion of Compx and Decompx programs in the October 19 and 21 release caused the replenishment system to fail on some of the AutoZone store servers. Replenishment system is a program used by AutoZone store servers to order and receive new inventory from the distribution centers.

Replenishment system failed with regard to the store service by distribution centers 66 and 77, approximately 650 of AutoZone's 3500 stores. And it says, investigation failed to reveal that the store servers served by these distribution centers still use the Compx and Decompx programs. In connection with the replenishment system, AutoZone IT department upgraded the software in distribution center 66 and 77 to eliminate the use of Compx and Decompx in all AutoZone stores. This upgrade occurred one day after the notice of the failure of the replenishment system. Correct?

A. Correct.

Q. You did this under your direction? Is that correct?

A. That's right.

Q. Now, am I correct that when you referred to the upgrade in 2003 of the replenishment system software, this is the replenishment system software you were talking about?

A. Yes.

Q. And you had originally believed that Compx and Decompx were not going to -- were not being used by any of these 3500 stores because of this upgrade.

A. That's right.

Q. And did you then learn at a later date that Distribution Center 66 and 77 had not been upgraded?

MR. STEWART: Asked and answered. You can answer.

MR. STONE: I don't think so.

MR. STEWART: You just asked that question. You just had him -- you just read this and asked him if that was right, and he said yes.

Q. (BY MR. STONE) All right. Just answer the question. Go ahead.

A. Yes.

Q. Okay. So they were using an older version of your replenishment system.

A. That's right.

Q. What is a distribution center?

A. Warehouse.

Q. So these two -- does this refer to two different warehouses, 66 and 77?
A. Yes.
Q. So these 650 stores got their parts from 66 and 77 warehouses, is that right, numbers 66 and 77?
A. That's right.
Q. Okay. So this replenishment system was the way that the stores would tell the warehouse we need more inventory. Is that right?
A. That's right.
Q. And then the warehouse would send the inventory to the store?
A. Right.
Q. So you learned that the -- up until this date that you upgraded the system, those 650 stores had continued to use Compx and Decompx to perform that process. Is that right?
A. Right.
Q. When you upgraded the system, did that -- strike that. Based on this upgrade that you made of the system, can you now represent on behalf of AutoZone that Compx and Decompx are not in any way being used by those 650 stores?
A. Yes.
Q. And you've deleted them from all of the 3500

stores servers?
A. Yes.
Q. Has Compx and Decompx been deleted to all the servers at AutoZone with the exception of the server where you're keeping your copies that you referred to earlier?
A. I think -- excuse me. I think that we are still using Compx and Decompx to transfer data between our mainframe and one of our AIX servers.
Q. And how long have you been doing that, to your knowledge?
A. Until -- let's see. I really don't know. It would have -- would have been a number of years.
Q. Would it have been that all the way back to the time that you were using OpenServer operating systems on your store servers?
A. I don't know.
Q. And do you intend to continue to use the Compx and Decompx for that purpose?
A. For some period of time, yes. But since Compx and Decompx aren't going to be universally used, it's in our best interest to standardize on something. And if this isn't it, we should probably get rid of it.
Q. But as you sit here today, you don't know whether Compx and Decompx contained SCO licensed code.

Is that right?
A. I don't know.
MR. STEWART: Is that for the application we just talked about, the continuing application?
MR. STONE: Yeah.
Q. (BY MR. STONE) Well, I mean, Compx and Decompx are programs. Right?
A. Yes.
Q. And is there some different version of the program that you're using?
A. Yes. Since we don't have source code, we buy, or bought, the product from the vendor. And we bought platform specific versions of that software.
Q. Who is the -- who is the vendor?
A. I want to say Compress Software, but I really don't know for sure.
Q. So you haven't made any -- taken any steps to determine who that vendor was prior to this deposition?
A. I didn't personally take any steps.
MR. STONE: We're going to request -- do you need to take a break?
VIDEO SPECIALIST: Yeah. We need to change tapes.
MR. STONE: Let's take a five-minute break.

VIDEO SPECIALIST: Off the record at 11:37. This is the end of Tape Number 2.
(SHORT BREAK)
VIDEO SPECIALIST: Back on the record at 11:49. This is the beginning of Tape Number 3.
MR. STONE: Well, off the record, since I guess that wasn't on the record, I requested the name of this vendor that you were referring to before we went off the record. And we'll write a letter, and Mr. Stewart will do whatever he does.
Q. (BY MR. STONE) Okay. So am I correct that that was an AIX operating system that this Compx and Decompx that you're using now runs on?
A. All on one side. MVS on the other, I think.
Q. So does -- can you represent to me that neither of those programs contain any SCO static libraries or SCO proprietary code?
A. I don't believe they do.
Q. If you move on to the third paragraph of Exhibit 36, the November 24th, 2004 letter from Mr. Stewart to me, it says the release script sent to the AutoZone stores on October 19th through 21 also searched each AutoZone store server for the existence of other COFF files beyond the files we identified to you previously. So these files that you're going on to

identify in this paragraph are files that weren't mentioned in any of the other discovery responses we've referred to. Is that correct?

A. Right.

Q. Okay. And it says -- you go on to say -- AutoZone goes on to say, a total of 1681 additional COFF files were found on 387 store machines. These programs had 127 unique file names. Correct?

A. Yes.

Q. And was that done under your supervision?

A. Yes.

Q. So there was a script that you sent out by your satellite. It asked questions of the machines and somehow gathered the information that there were these additional COFF files on 387 of these 3500 machines. Is that right?

A. That's right.

Q. Did you do anything to try to determine why it would be that 387 store machines had these files on them but the other store machines did not?

A. We didn't do anything in particular to figure out why these files were on these stores. Because if I remember correctly, none of these files were in the places where production executables ought to be, or the names had been changed.

For example -- let me think about this. Historically, if a store had a program -- problem with a program, one of the actions that the help desk might take would be to verify the checks on the executables. And if they found that that executable was altered in some way, they would -- they might rename the file to be dot bad and then send them a new copy, a correct copy of the file, of the program as a way to fix the problem with the program.

A number of these files had names like, you know, dot Bob, dot Joe, dot Steve, which is indicative of one of these situations where a copy was made for some reason, and then a fresh copy would typically have been sent to the store, also.

Q. So those were probably -- you believe based on what you've seen and your knowledge of how things work at AutoZone that the reason that these files were on these 387 store machines were that they were put there by this support team in connection with some work that they were doing or question they were answering?

A. Support team or developers. Right.

Q. When you say 1681 files on 387 store machines, does that mean that a total of 1681? You don't mean that there's 1681 on each machine. Is that right?

A. No. There's a total of 1681 files sprinkled across 387 machines.

Q. And of those 1681 files, they have -- 127 of them have -- I don't know exactly how to say this. But there's only 127 actual different file names?

A. Right.

Q. But some of the files are actually different even though they have the same file name? Is that right?

A. No.

Q. There's many copies of the same file?

A. Yes.

Q. Oh, okay. And you deleted all those files again using this satellite system on November 1st? Is that right?

A. Yes.

Q. Or by November 1st?

A. That's right.

Q. So these were additional files that may have contained SCO licensed codes that were, in fact, copied on to the Linux OpenServer systems at AutoZone. Correct?

A. Right.

Q. Then you go on in the next paragraph, or AutoZone goes on in the next paragraph to say that you discovered that a machine load computer had a program on

it called dexpand.x -- dot X, that was compiled under OpenServer. How did that happen? How did you find that?

A. As part of our release process, once a release has been sent to the majority of the chain, we're able to put that same release on the computer that we used to create new store loads so that before we open a new store, when we create an image for that store, it gets created with the latest software.

So as is standard procedure, once I had sent this release to essentially the chain, I also sent this release to the store load computer to find and remove any COFF files that might be on that machine. And this dexpand program was on the load machine, and the release did remove it.

Q. Could you tell whether that expand program had been used during the time that the -- or after the migration to Linux?

A. I believe it was used after the migration to Linux.

Q. And do you know whether that expand program had SCO libraries in it?

A. I believe it was a statically-linked file. So it probably did have SCO code in it.

Q. And I think you probably said this, but what

did the program actually do?
A. It converted data formats from a pre-Y2K format, that was in use on the mainframe, to a post-Y2K format that is required by the stores.
Q. And that was -- was that something that had to be done every time a program that had that pre-Y2K format was being sent to the stores?
A. The process of creating an image for a new store involves loading that image with programs but also with some of the data that the store is going to need.
So when we decide today's the day to build an image, part of that process involves running some programs on the mainframe that generate data. We move that data to the store image machine. We convert its format along the way. And that way when the computer arrives at the store, it already has some foundation data to go along with the executables.
So that program was used basically once every time you would open -- every time we wanted to open a new store.
Q. Okay. So every time you opened a new store. Then you -- how did you reco -- you say that you recompiled this program. What did you do to recompile this program?
A. I didn't personally recompile it. So I don't

know whether there were any code changes necessary. But I do know it only took a short period of time to recompile and test. So the changes could not have been extensive.
Q. And as you sit here today, you can testify that the current version of this program does not contain any SCO licensed code?
A. That's right. It was recompiled on the Linux machine for Linux.
Q. And then what did do you? Did you take a previous copy and make a copy and put it on the Spirit machine?
A. Vision.
Q. I'm sorry. The Vision machine?
A. (Witness nod affirmatively.)
Q. Is that what you did with all the -- anything that you deleted, any program that you deleted?
A. Yes.
Q. And would that copy contain the information about, you know, when the program was deleted and when it was loaded on the machine?
A. I don't know. For some of the copies of COFF files over to Vision, the system administrators used a utility called CPIO, which may retain the original date on the file. So that when you load it on to Vision, it

will look as if it's been there for a long time.
For other files, that's not the case. We FTPed a copy from -- used a utility called FTP to transfer a file from wherever it was originally on to Vision. And in that case, it would retain the date on which it was copied.
Q. Okay. Was dexpand.x a binary that was created with the SCO developer kit to run on OpenServer?
A. Originally, yes.
Q. By AutoZone?
A. Yes.
Q. Okay. Go on, it says AutoZone's IT personnel also discovered that the Spirit server had some OpenServer compiled programs on it because of a recent restoration of the server after it crashed. The Spirit server is -- isn't that an OpenServer machine?
A. Spirit was an OpenServer machine. We were having some problems with Spirit. In particular, we knew the disks were going bad. And so our system administrators were preparing a replacement machine. Well, apparently, they didn't prepare it quickly enough. Because one evening, Spirit died. And so we worked through most of the night trying to restore all of Spirit's data on to that replacement machine.
Q. The replacement machine that was being

prepared?
A. Right.
Q. Is that what you mean?
A. Right.
Q. And in doing that, you loaded 1130 SCO compiled programs on to the machine? Is that right?
A. Right. We restored everything onto that replacement machine, and that included some SCO binaries.
Q. And what was that replacement machine going to be used for after that?
A. It was going to continue to be used for our software repository.
Q. Okay. So this is this revision control system that you had spoken of earlier where you keep your source code? Is that right?
A. That's right.
Q. So had you originally intended to remove the SCO code from that machine and only have AutoZone's home grown codes, so to speak?
A. Right. In an orderly migration, we would have copied the repository, and then we would have copied only the other directories that we needed rather than just a wholesale replacement of everything.
Q. So based on what you've now done, did you do

something to remove all the SCO licensed code from that machine?
A. Yes.
Q. So that was one machine that these 1130 programs were on. Right? It was not --
A. Right.
Q. Do you know what -- what operating system did that machine use?
A. It was OpenServer 505, I think, before the crash, and then it was Red Hat. I don't -- I really don't know what it is right now. I believe it's 7 dot 2, but it might be Red Hat Version 9. I really don't know.
Q. So it's some version of Red Hat's Linux distribution.
A. That's correct.
Q. You switched it from a OpenServer machine to a Linux machine after the crash.
A. That's right.
Q. But when you did that, you copied these 1130 programs on to the machine?
A. Right.
Q. When did that occur?
A. I'm sorry. I don't remember the date.
Q. Can you give me the year?

A. No. I really can't. I don't know when that was. I believe it was last year, but I really don't know.
Q. So there may be a year that you were using this machine where it had the 1130 -- the 1,130 SCO files on it. Is that right?
A. I don't know how long that Linux machine was up before we removed these files.
Q. Do you know whether any of these files were used during that time period?
A. I don't think any of these files would have been used because they would have been in individuals' home directories. And that machine is no longer used as the machine that people log into to perform work. So...
Q. What machine is used as the machine that people log in to perform work?
A. Only a machine called Wrangler and a machine called Jeep for Linux development.
Q. So Jeep and Wrangler are two different machines?
A. (Witness nods affirmatively.)
Q. Is this something to do with the Jeep Wrangler or just a coincidence?
A. Our system administrators find it an amusing policy to name machines after cars.

Q. So Wrangler was the machine that you originally used, right, to develop Linux?
A. Yes.
Q. Did Jeep come after that or something?
A. Jeep came after that.
Q. Did you look on Jeep to see if Jeep had any COFF files on it?
A. Yes.
Q. Did you find any?
A. I think we did, but I really don't remember.
Q. But they would be in that -- the CDs or that hard drive that you produced to us. Is that right?
A. Yes.
Q. It goes on and says, finally, we've reviewed the relevant OpenServer agreements between SCO and AutoZone. These agreements are still in place and do not include any prohibitions on AutoZone's use of OpenServer compiled code on Linux machines. Did you have anything to do with that review?
A. I tried to find copies of these agreements. But no, I didn't have anything to do with the review.
Q. Did you find any copies of the agreements?
A. I did find some documents that I forwarded to the attorneys.
Q. Can you tell me the names of the documents?

Do you remember?
A. No. I'm sorry. I don't remember.
MR. STONE: Okay. This seems like a good time to break for lunch. Actually, can we stay on for one second?
Q. (BY MR. STONE) I just have one question about something that we had gone over earlier. It's on Page 2 of Exhibit 36. It says, In addition to the -- it's in the first paragraph like five lines down -- five lines down. In addition to the help utilities, there were four SCO compiled programs that were copied to the Linux machines errantly because the programs were not located in the proper directory. We talked about the fact that that happened. Correct?
A. Yes.
Q. I just wanted to know what these four programs were intended to do, if you remember.
A. I'm sorry. I don't remember. I believe they were located in the directory where we usually keep data. But I don't -- I don't remember what the programs were.
Q. You don't remember what they were actually intended to do? Is that right?
A. That's right.

Q. Okay.
MR. STONE: Let's take a break for lunch.
VIDEO SPECIALIST: Off the record at 12:10.
(LUNCH BREAK)
VIDEO SPECIALIST: Back on the record at 1303.
MR. STONE: Would you mark this as the next exhibit, please.
(WHEREUPON, THE ABOVE-MENTIONED DOCUMENT WAS MARKED AS EXHIBIT NO. 37, AND IS HERETO ATTACHED.)
Q. (BY MR. STONE) Mr. Celmer, I'm showing you a letter dated March 4, 2005 from your counsel to counsel for SCO. And have you seen that before?
A. I think so.
Q. Now, were you -- were you the person responsible for putting together the data that was produced to us with this letter?
A. I did part of it, and I oversaw part of it.
Q. Okay. Can you tell me what are the -- the denominations FEP1 through 8, what do those refer to?
A. That refers to a set of machines that we use presently and historically to communicate with the stores. Our stores are updated daily with data, and of

course, they send back sales information and so forth. And these are the machines that we use to pass that data back and forth.
Q. And did you find some COFF files on these machines? Is that why you produced images from them?
A. Yes. I think so.
Q. Do you recall whether there's any document which would tell us which -- what COFF files you found on those machines, or have we referred to it already?
A. I'm not familiar with a document that describes what was found on these machines.
Q. Okay. But whatever you found would be included in what you produced to us. Is that correct?
A. That's right.
Q. How about the VEGA server? What is the VEGA server?
A. I think VEGA is a print server. And evidently, we found one or more COFF files on that machine, as well. VEGA, I believe, used to be an SCO Unix server. So that may be part of the explanation for how a COFF file got there.
Q. And was it converted later on to a Linux server?
A. I think so. Yes.
Q. Do you know when?

A. No. I'm sorry.
Q. Where -- is the VEGA server located in your headquarters?
A. Yes, it is.
Q. What is it -- what is it for?
A. It is a print server. I think that our design group uses it throughout their printing.
Q. Then if you go to Page 2, Spirit, we talked about. Correct?
A. Yes.
Q. So this was an image of the Spirit server that we've been referring to which -- I've gotten these confused. One of them was a source code repository, and the other one a Vision or a development machine? Is that right? Which was Spirit?
A. Spirit was a source code repository.
Q. And Spirit is the one that crashed and that you had to create a new machine for. Is that right?
A. That's right.
Q. And then you gave us the image for the Jeep server?
A. That's right.
Q. And that is another development machine? Is that right?
A. That's right.

Q. And is that a Linux development machine?
A. Yes, it is.
Q. And you found COFF files on that, as well?
A. Yes.
Q. Do you recall how many you found?
A. No, I don't. But when Jeep was commissioned, so to speak, as the second Linux development machine, some of the users of Jeep had their directories mirrored between Jeep and Wrangler. So some of the same files that were on Jeep -- I'm sorry -- Wrangler, were copied to Jeep as a result of that mirroring.
Q. So if we go to the next item on the list, Wrangler, Wrangler was the original Linux development machine. Correct?
A. That's right.
Q. You still use it now as a development machine?
A. Yes.
Q. And that contained COFF files on it which you're saying when they created the Jeep development, Linux development machine, were transferred or copied to that machine as well? Is that right?
A. Right.
Q. Do you remember how many files were on either the Jeep or the Wrangler machines?
A. No. I'm sorry. I don't remember.

Q. Okay. But they would be in addition to the files that we've discussed to this point in the letters that I've shown you and discovery responses. Is that right?
A. Right.
Q. Then it says image of store server releases. These are releases that were made through your satellite to the store servers? Is that right?
A. I believe that at one point, we created for you a list of every software release, not just a list, but a copy of every software release that we still had that we'd ever sent to a store. And I think that's what this is. So really, that's store release server.
Q. Oh, I see.
A. So the server from which we send out releases, that's what we gave you a copy of.
Q. Okay. And you found that there were COFF files in those releases. Correct?
A. Right.
Q. In some of them?
A. Right. Originally the stores were SCO Unix. And so all of our releases contained SCO COFF binaries.
Q. But as well, even afterwards when you switched over to Linux, some of those releases contained COFF files, as well. Correct?

A. Some of them did. Certainly during the transition, we had to support software on both formats. And then subsequent to the transition to Linux, I don't -- I don't think we created any more COFF binaries. There wouldn't have been a need to.
Q. Didn't you tell me that there's at least one release in which you did release several COFF files? Whether intentionally or unintentionally, they were released in one of the releases. Correct?
A. Right. That was during the transition period while we had both SCO and Linux with the stores.
Q. Okay. And then there are four store servers, is that correct, images?
A. The -- this document refers to four store images. Right. You guys had requested five.
Q. Yeah. That's what I thought.
MR. STONE: Do you know, have we only received four? My associate's telling me that one of your servers crashed or something and we only got four instead of five? Is that right?
MR.BRIDGES: Yes.
Q. (BY MR. STONE) So these would be images of the store servers as they currently exist or as they existed at the time you made the image which was after you deleted all the files we've discussed?

Is that right?
A. That's right.
Q. So if we looked at those store server images, we should not find any files that contain SCO proprietary code or licensed code. Is that correct?
MR. STEWART: Object to form. You can answer.
A. I don't think you'd find any COFF files.
Q. (BY MR. STONE) We might find some Zenix files?
A. Right.
Q. This is the image of the store servers that you're currently using in your stores as far as you know? Is that right?
A. Which one is the image?
Q. Well, I thought the idea was that your stores essentially all use the same image of -- for the Linux operating system that's in place at that time. Is that right?
A. That's right.
Q. So that these should be representative of the other I don't remember how to -- if you took 3500 and then you took four away from it are stores that are running Linux. Is that right?
A. They would look very much like this.

Q. Okay.
A. Minor differences.
Q. And then you gave us copies of the files that you identified using the script that we provided to you. Is that correct?
A. Right.
Q. You understand that our expert provided us with a script to help you locate additional files and you agreed to run that. Correct?
A. Yes.
Q. And you did run that, and it identified some additional files. Is that right?
A. That's right.
Q. And was that 15 ELF and Zenix files, SCO, ELF and Zenix files? Does that sound correct to you?
A. I don't know how many files that that program identified. You know, we produced a list of the output for every store in the chain. And I believe I summarized it just picking out the ones that were of interest. But I don't -- I really don't remember the numbers.
Q. Okay. So this last bullet on Page 2 of the exhibit output of SCO's search program run on every AutoZone store server, that's what you're talking about?
A. Yes.

Q. So if we look at that, that should show us what files it created. Is that right?
A. Yes.
Q. Or it found. Okay. So I just want to go over this to make sure I've got everything and that there's nothing else you know about. Okay?
A. Uh-huh (affirmative response).
Q. There is 1681 COFF files that you found on 387 AutoZone store machines. Is that correct?
A. Yes.
Q. There's 19 unused COFF files that were loaded onto the Linux image and were for some period of time on 3500 of the SCO machines. Correct?
A. Yes.
Q. There is nine COFF files that you don't know whether they were used or not used that were loaded on to the Linux image and for some period of time were on the 3500 machines. Correct?
A. Yes.
Q. There is two COFF files, Compx and Decompx, that were on the 3500 machines. Correct?
A. Yes.
Q. There's dexpand dot X, which was on your machine load computer which was one machine in the headquarters. Is that right?

A. That's right.
Q. There's 1130 SCO compiled programs that were on your Spirit server. Correct?
A. I don't remember how many were on Spirit.
Q. I'll represent to you that was the number that was given to us.
A. I believe that number.
Q. Okay. There's -- and you don't remember whether it's 15 or some other number of SCO, ELF, and Zenix files. Is that correct? I just asked you about those.
A. I don't know how many ELF and Zenix files we found.
Q. Okay.
A. SCO ELF, rather.
Q. Then in addition to that, there would be whatever would be on FEP machines that you've produced to us. Right?
A. That's right.
Q. Whatever would be on the VEGA machine. Correct?
A. That's right.
Q. Whatever would be on the Spirit, Jeep and Wrangler machines. Correct?
A. Right.

Q. Whatever would be in the store releases.
A. Right.
Q. And whatever that script found. Is that correct, the script that we provided to you?
A. Right. We talked about that one.
Q. Is there any -- all right. One of the -- one of my colleagues believes that I may have misspoken and said loaded onto the SCO machines. I want to be clear that all of the files that I've been talking about were loaded on AutoZone machines. Is that correct?
A. That's correct.
Q. And are you aware -- other than what I've just named, are you aware of any other files on any other AutoZone machines that are COFF files or Zenix files?
A. No.
Q. Has, to your knowledge, AutoZone disclosed to SCO all the files that were at one time copied on to any of its Linux operating system machines?
A. I'm sorry. Say that again.
MR. STONE: Could you read it back?
(WHEREUPON, THE LAST QUESTION WAS READ BACK BY THE COURT REPORTER.)
A. Yes.
Q. (BY MR. STONE) All right. I want to go back and ask you a few questions about the migration

process to make sure I understand how it -- how it worked. Okay? You started out with binaries that were compiled to work on OpenServer. Correct?
A. Right.
Q. To operate on OpenServer. And you wanted to go to binaries that would run on Linux. Correct?
A. Right.
Q. Did you do anything to test or compare the binaries that you were creating to run on Linux with the binaries that were running on OpenServer to make sure that they were creating -- generating the same output?
A. No. We tested the functionality of the programs, but we didn't compare the files.
Q. What do you mean by saying you've tested the functionality of the program?
A. If given a program that performed some function on OpenServer, after recompiling that program on Linux, we would execute that program and verify that the functionality under Linux is the same as the functionality under OpenServer.
Q. Did do you that by running the open -- the program on -- that was written for OpenServer and compare the results of the two?
A. Yes.
Q. Is that what the -- I know I'm going to get

this wrong -- Vision machine was used for?
A. Most of our programs could be run on and tested on Vision being in the development box. Some programs would have needed the fuller environment provided by our test machines, which are more like store computers.
Q. Like, to replicate the store environment?
A. Right. So most of the time, you can run a program on Vision, find out how it works, take a look at output whether it's on a screen or in a file, and then go to a Linux machine to run the Linux version of that program and verify that the screens look the same and the file output's the same.
Q. That was the question I was asking. So is that something that you did when you were porting these binaries to -- I'm sure I'm using it wrong again. When you were -- when you were doing the port process from OpenServer to Linux, you would do whatever you did to create your binaries to run on Linux, and then you would compare that with how the original binaries ran on OpenServer and compare the output?
A. Right.
Q. Did you need to be able to run those binaries on the same machine to do that, or did you need to be able to read the output on the same machine to do that?

A. No. We almost never work on the actual machine. We work on our machines at our desk, and we simply communicate with these those development machines. So typically, we would have one session which is either on Vision, the OpenServer development machine, or it's connected to one of our OpenServer test machines. And then on a separate machine, we would be connected with Wrangler, the Linux development machine, or else one of the Linux test store machines.
And so we would either alternate between them. Or if you were sufficiently familiar with the program, you might not have to refer to the OpenServer version at all because you know what the output's supposed to look like.
Q. So you had both these machines up and running, and then you'd call upon them depending on what it is that you needed from your own machine. Is that right?
A. Right.
Q. And is that one of the ways that you would determine, for example, what flags needed to be put in the source code or what modifications might be needed to make -- be made to the source code so that the source code would compile on both OpenServer and Linux?
A. Sometimes. Sometimes you would find differences in the behavior of the programs based on

simply running it and noticing that it's different. You can also, upon finding one of those differences, record it in a list of things to be changed when you do the next port. So if you find that in order to execute some command on OpenServer, it has a certain syntax. But if you want to run the same command on Linus and it has a different syntax, that you add it the list of things to check for in your porting. You look for that first syntax. If you find it, you immediately change it. There's no need to find that same error a second time.
Q. When you were doing this, would you refer to the -- to the manuals or the other documents that you had for OpenServer to understand how the programs were supposed to run?
MR. STEWART: Object to form. You can answer.
A. I suppose if you encountered a command and you didn't know what one of the options for that command was, you might read the man page on the original OpenServer platform and find out what it was doing. And similarly, if you didn't know already the syntax under Linux, you would consult the manual page for the Linux version to find out how to make it do what you wanted it to do.
Q. (BY MR. STONE) So what did you do? Did

you take all these binaries and sort of divide them up and say, Developer A, you take these binaries and you make binaries that will work on Linux, and Developer B, you take these, and you just sort of parcelled them out to different people?
A. Mostly, that is how it would work. You'd find a collection of programs that work together, and you'd ask someone to change all of them so that you can run all of them and test the inner activity.
Q. Did AutoZone ever share its software with any of the other stores that it owned or acquired other than the AutoZone stores?
A. There really are only a couple of meaningful subsidiaries to AutoZone, one being a retail truck parts company, and the other being a provider of software to repair shops, software and data to repair shops. The only sharing I can think of right now is there was an effort to make our parts lookup system available to the truck parts company. So we -- I really don't know how much software modification was necessary. But we had to try to make the existing program available for that -- for that company.
Q. Do you know whether -- did you do anything to determine whether any of those programs that you made available to the truck subsidiary contained COFF files?

A. I don't -- I don't know whether anyone looked to see whether there were COFF files on that machine.
Q. Okay. Where would that -- where would that machine be located?
A. AutoZone divested itself of that company a couple of years ago.
Q. Do you know what happened to the machine?
A. I really don't know what happened to the machine.
Q. What about -- is there any way for you to determine what would have been on that machine from information within Auto -- currently within AutoZone?
A. I don't know. I don't know the extent to which the hardware is traceable by asset tag, for example. And I also don't know whether we allowed that company to continue to have access to all parts lookup software. So it may have been that we didn't want them to have that software. So we kept it and the machine. But I don't know.
Q. How about customers of AutoZone? Did you ever share any of this software that you used in the stores or that you used in your headquarters with customers?
A. We talked about allowing some of our commercial customers to have access to our parts lookup so that they could look up AutoZone parts and purchase

them from us. I don't know whether we ever installed one of our computers in any customer location. No.
Q. If -- would your parts lookup software have contained COFF files? Do you know?
A. It would have depended on the timing. If it was, you know, while we were using OpenServer, then it would have been OpenServer software.
Q. But what if you were to have just, you know, given them programs off your Linux image which contained COFF files? Is it possible they could have gotten them that way?
A. If we just created their image from our store load image, then it is possible that they would have had COFF files on those machines.
Q. How could we determine whether that occurred?
A. You would have to find out from our commercial business executives whether we had ever put one of our computers in a commercial customer location.
Q. Okay.
MR. STONE: I'm going to make a request for that information. I'll follow it up with a letter, and you can respond however you respond.
Q. (BY MR. STONE) Just to clarify something that I asked you earlier, there were occasions which you would -- in which you would compare the output

that you would get from the Linux binary that you compiled to work on Linux with the output that you would get from the OpenServer binary that AutoZone had compiled to run on OpenServer. Correct?
A. I'm sorry. Say that again.
MR. STONE: Would you read that back?
(WHEREUPON, THE LAST QUESTION WAS READ BACK BY THE COURT REPORTER.)
A. That's correct.
Q. (BY MR. STONE) And if there was something different in the output, that would lead you to make some change to the binary that you had compiled to work on Linux, so it would create identical output if possible. Is that correct?
A. That's right. We almost always wanted to exactly duplicate the functionality under Linux that we already had under OpenServer.
Q. Was the first implementation that you did of Linux, you did that through Red Hat? Is that right?
A. The first store implementation. Yes.
Q. And you believe that was some time in the 2000 time frame? Is that right?
A. I think we started porting the software in '99 or 2000. As I said, we could run most of those programs on the development machine. So the creation of

the first -- the creation of the first Linux store machine even for test purposes probably was 2000.
Q. Was that -- do you recall, was that Linux Red Hat 6.2?
A. Yes, it was.
Q. And did you confer with anybody else before -- before you decided to use Red Hat Linux 6.2 and anybody else who had used it to see what their experience was with it?
A. Greer was the one who originally ported our AutoZone libraries and first couple of critical applications just to see whether it could be done and how it worked. He also was responsible for investigating which version of Linux we were going to try. And so he's the guy that settled on Red Hat, and 6.2 was a version that was available when he began. And so it was the version that was stable by the time he was ready to roll it out.
Q. Did you have any discussions with anybody at IBM about implementing Red Hat 6.2?
A. I don't think so.
Q. Did they play any role in your implementation of Red Hat 6.2?
A. Our implementation of -- no.
Q. When you implemented Red Hat 6.2, to your

knowledge, did anybody that you were working with review SCO licenses in connection with that implementation?

A. I don't know of anybody that reviewed licenses for the Linux rollout.

Q. Did there come a time when you implemented a different version of Red Hat Linux?

A. Yes.

Q. When approximately did that happen?

A. In -- I don't remember -- 2001 or 2002, Greer left AutoZone. When he did, his effort to port our software to Linux was largely successful. We believe it to be largely successful. So we were going to go ahead and roll Linux out to our stores.

One of the things that concerned me most -- and I took over the -- I took over the responsibility of rolling out Linux. One of the things that concerned me most in Red Hat 6.2 was that they were using a file system that was not journaled. We had a lot of stores, and they are in sometimes areas where the power is questionable. And anyway, having a store just sort of crash without a journaling file system was something that I didn't really want to go forward with.

By the time we were ready to go forward, Red Hat 7.2 was out, available, and stable in our opinion. The EXT3 file system was available for that operating

system. So that's what we chose to do. We chose to -- instead of go forward with Red Hat 6.2, we chose to go forward with 7.2 because of the availability of the journal file system.

Q. So that was in the 2000 time frame, or when was that?

A. '01, '02, something like that.

Q. In the 2001 to 2002 time frame?

A. I think so.

Q. That's when you rolled it out to the 3500 stores?

A. That's right.

Q. Did there come a time period after that that you implemented, yet, a newer version of Red Hat Linux?

A. Not in our stores. There are a number of fixes that have become available since the time we rolled that out, and there's some increased functionality, as well. And so we're always weighing the benefits of upgrading to a newer and presumably better operating system versus the cost of doing that and the risk of doing that.

And so we put Red Hat Version 9, I believe, on these FEP machines when they were migrated from OpenServer to Linux. And part of the reason for doing that was to give us some familiarity with -- with Red

Hat Version 9 to decide whether we could use it in the stores or should use it in the stores.

Q. So the stores currently are running Red Hat 7.2?

A. Still, yes.

Q. And am I correct that you believe at some time in 1999 that you first started porting programs to Linux?

A. I think that's right.

Q. And that the rollout, you believe, occurred -- do you believe it occurred some time in 2001, that it started, or you just don't recall, as you sit here today?

A. I really don't remember the dates.

Q. So I take it you don't recall when it ended, either. Is that correct?

A. Sorry.

Q. I'll just ask this question: What JFS do you use in Red Hat 7.2?

A. We use EXT3 file systems.

Q. Does that -- does JFS mean journal filing system?

A. I believe it does.

Q. Do you know what an interprocess communication capability is?

A. Yes. I think so.

Q. Can you tell me what it is?

A. Interprocess communication capability is support for allowing one program to communicate in some way with another program.

Q. Did that exist in OpenServer?

A. Yes, it did.

Q. Do you know, does that exist in Linux?

A. I believe it does.

Q. As you sit here today, could you tell me the AutoZone programs that utilized this IPC capability?

A. I can't tell you all of them. And we would have to agree on the definition of IPC. Because to some extent, any program that writes to a file can be communicating with any other program that may read from that same file.

Q. Well, I'll agree that whatever you say IPC is is what I say it is, because I don't know what it is. I'll just accept your definition. So with your definition, could you tell me which AutoZone programs?

A. With that definition, there would be a lot of communication between programs. So, no, I couldn't.

Q. You would define it broadly?

A. If we were to decide to define it broadly, then there would be a lot of communication going on.

Q. Do you know where the IPC capability in Linux comes from?
A. Well, again, we sort of have a problem with definition. In using the same broad definition, it would come from the files.
Q. Are you familiar with something called -- that's SCO proprietary product called System 5, System 5 Version 3 and System 5 Version 4?
A. I am familiar with the System 5 Unix.
Q. And that's a version of Unix. Correct?
A. I believe it is.
Q. Do you know if the IPC capability in Linux comes from System 5?
A. I don't know.
Q. Do you know whether the IPC capability in OpenServer comes from System 5?
MR. STEWART: I'm going to object to this continuing down this road unless I see that this is tying back into the conversion. Because issues related to infringement in Linux of Unix has been stayed at this point.
MR. STONE: Okay. You're saying you object --
MR. STEWART: Can you tie this into the conversion?

MR. STONE: I have like -- I don't have many questions on this, and I doubt I'm going to get any answers. So but --
MR. STEWART: I mean, I'll let it go a little bit longer, then.
MR. STONE: -- I understand your objection.
Q. (BY MR. STONE) Do you have an answer to that question?
THE WITNESS: Can you read it back, please?
(WHEREUPON, THE LAST QUESTION WAS READ BACK BY THE COURT REPORTER.)
A. No, I don't.
Q. (BY MR. STONE) Okay. Just to clarify something I asked you earlier, have you heard of Car Max and Tire Kingdom?
A. Yes.
Q. Are those commercial customers of AutoZone?
A. Yes.
Q. Do you know whether AutoZone installed Linux in Car Max and Tire Kingdom servers?
A. I believe we installed Linux in the Car Max locations. Car Max locations are essentially an AutoZone store with only one customer, Car Max. I

believe that AutoZone employees run that system, for example.
Tire Kingdom, I believe, is using a server inside AutoZone headquarters. Tire King -- or was using. I'm not sure they're even a customer anymore. I think all they were doing was parts lookup on that server. And Car Max locations are a little different. They are -- since they are essentially an AutoZone store, they use most of the suite of applications. But because they only have only one customer, they don't necessarily use all of them.
Q. Now, are these Car Max locations included in the 3500 AutoZone stores that you referred to earlier, or are they in addition to those 3500?
A. They're included.
Q. Do you consider them part of the 3500 --
A. Yes.
Q. -- AutoZone stores? And as to Tire Kingdom, you don't know whether Tire Kingdom would have been having access to machines containing COFF files as you sit here today? Is that correct?
A. I don't think they had access to the machines. Well, I think they only had access to the machines as users, but it would be a least line at AutoZone.
Q. What about stores in other countries like

Mexico? Does AutoZone have stores in Mexico?
A. Yes.
Q. Is that included in the 3500 stores that you've referred to?
A. It may or may not have been included in the store count. We only have 60 or so stores in Mexico. So I really don't know for any given count whether it was included.
Q. Did you look on those stores in Mexico to make sure they didn't have COFF files on them?
A. Yes.
Q. Do you recall whether you found any COFF files on those stores?
A. We did find COFF files for those stores.
Q. And you -- and you deleted them or --
A. Yes.
Q. -- recompiled them?
A. Right. Treated them just like any other AutoZone store.
MR. STONE: Let's go off the record for a minute.
VIDEO SPECIALIST: Off the record 1351. This is the end of Tape 3.
(SHORT BREAK)
VIDEO SPECIALIST: Back on the record.

The time is approximately 1412. This is the beginning of Tape Number 4.

Q. (BY MR. STONE) Mr. Celmer, just a couple of questions off the topic of what I was on before. Did you speak to anybody other than your counsel before appearing here for your deposition about the deposition?

A. No.

Q. Did you speak to Mr. Greer at all?

A. No.

Q. Did you meet with your counsel prior to appearing here today for your deposition?

A. Yes.

Q. When did that occur?

A. Just a few minutes before the deposition.

Q. You didn't meet with them on any other previous occasion to prepare for this deposition?

A. We had a conference call on Friday, Thursday. Thursday, I think.

Q. And how long did that last?

A. Couple hours, I suppose.

Q. Was there anybody other than Mr. Stewart and Mr. Bridges on the call, and you?

A. Becky was on temporarily.

Q. That's AutoZone's general counsel? Is that

right?

MS. BALLOU: Assistant.

MR. STONE: Assistant. Well, I'm trying to promote you.

Q. (BY MR. STONE) Did you review any documents before coming here for this deposition?

A. I think I looked at the deposition notice, wherever it is.

Q. Of Jim Greer?

A. No. The notice for this deposition.

Q. Oh, okay. The 30(B)6 notice --

A. Right.

Q. -- that I had marked?

A. Right. That's the one.

Q. Did you look at Mr. Greer's transcript?

A. No.

Q. Are you aware of a web posting that Mr. Greer made on a website called Groklaw?

A. I've heard rumors of such a posting.

Q. Do you know how that came about?

A. No. I really don't.

Q. Did you ever speak to him about it?

A. We have -- we, you know, sort of joked, slash, commiserated about the suit. We've never really talked that much about it.

Q. Did you talk to him after he had his deposition?

A. I'm sure I spoke to him after he had his deposition. But I didn't really -- we didn't really talk about the deposition itself.

Q. Have you talked with Mr. Greer about this case in any substance other than, you know, that it's going on and commiserating about it?

A. Not really.

Q. Did you review that posting in preparation for your deposition?

A. I don't think I've seen that posting.

Q. Getting back to some testimony that you gave earlier, you were testifying about why in your opinion it was preferable to use Linux 7.2 as opposed to 6.2? Correct?

A. Right.

Q. And you mentioned that 6.2 did not have a journal file system? Is that right?

A. As I recall, we either couldn't use EXT3 at all under 6.2, or we would have had to apply some patches in order to support it. And rather than do either of those things, I thought it best to simply upgrade the whole operating system and then simply select EXT3 as an option.

Q. And EXT3 is -- what is EXT3?

A. I believe it's just an extension of the file system such that if a computer loses power, for example, when the machine powers up, the file system structure itself is preserved and doesn't have to be cleaned up. You may lose data, but at least the arrangement of files on the disk is clean.

Q. So is this a program or a series of programs?

A. Well, yes. Effectively, everything on the computer is a program. And the part of the operating system that manages your files for you is what I'm referring to when I talk about EXT2, EXT3.

Q. And I think you mentioned before, you were concerned because you thought some of the AutoZone stores were in locations where they might lose power.

A. Definitely?

Q. Is it correct that OpenServer 5.05 had a journal filing system?

A. It had a number of file system options, one or two of which were journaled, yes.

Q. So you had that capability when you were using OpenServer. But when you moved to 6.2, you didn't have it. Is that right?

A. In the 6.2 -- four 6.2 test stores, we were not using a journal file system, and I think that's

because it wasn't -- a reliable version wasn't available at that time.

Q. Do you know where ETX.3 comes from?

A. No, I don't.

Q. Okay. We were also earlier talking about your recom -- what AutoZone did to recompile programs that were in COFF format that you found that you still wanted to use. Do you recall that?

A. Yes.

Q. In the -- during the discovery process.

A. Yes.

Q. And I think you told me that you recompiled them in the ELF format? Is that right?

A. Right.

Q. Is that the format that you recompiled all of the programs that you ported from OpenServer to Linux when you first did the migration in the 2000 time frame?

A. I think so. Yes.

Q. And is that the format that all the programs on your 3500 machines that are running Linux operating systems are in now?

A. I think so. Yes.

Q. Do you know if AutoZone did anything to determine who owned that format or if somebody had proprietary rights in that format?

A. I'm not aware of anybody investigating that.

Q. I'm going back and asking follow-ups on things that you've said. One of the things you talked about was having a machine and looking at the -- was it the Vision server -- speaking to the Vision server and speaking to Wrangler server? Is that right?

A. For example, yes.

Q. So Vision would run on OpenServer, and Wrangler would run on Linux. Correct?

A. Right.

Q. And you mentioned that you would get output on occasion from a -- from the Linux machine, and you'd want to compare it with that same -- the equivalent binary running on OpenServer. Correct?

A. Right.

Q. Did that output, was that ever in a binary form?

A. Yes.

Q. So did you write scripts or something to compare the binary outputs between the two?

A. The binary output I'm talking about is largely data, data files. They're just not human readable. For each of those files, we also have utilities that are used to translate the contents of those files into human readable forms.

And so what we would typically do is run the programs to create the data in the files but then use the utility to display the contents of those files for us so we could read them.

Q. And then you'd just compare them to make sure that the output was equivalent.

A. Right.

Q. So you didn't have to write a script that would like compare line by line or something like the output was. Is that right?

A. That's right.

Q. All right. I'm going to say a number of things. I'm not sure I know what any of them mean. I'll ask the question, and you can try to answer it. Okay?

A. Uh-huh. Yes.

Q. I asked you some questions before about IPC. Do you recall that?

A. Yes.

Q. If I were to devine -- define IPC as Semaphore -- Semaphore's message queues and shared memory --

A. Okay.

Q. -- could you tell me what AutoZone programs would utilize IPC?

MR. STEWART: I'm going to object to the extent this is outside the deposition notice. I'll allow you to answer it.

MR. STONE: Thank you.

MR. STEWART: I hope we're not going to go too far down this road.

MR. STONE: We're not going to go too far.

A. I can certainly tell you some of the programs that used those mechanisms. I'm not sure I can give you an exhausted list.

Q. (BY MR. STONE) Let me try this: First, I'll read some names of programs, and you can tell me whether it does. Okay?

A. All right.

Q. MQ Series?

MR. STEWART: Same objection. Could we have a standing objection?

MR. STONE: Yes.

A. I believe MQ series used Semaphores and shared memory. I don't think MQ series used message queues, but I'm not certain about that.

Q. (BY MR. STONE) Okay. PayCalc?

A. Say that again.

Q. PayCalc?

A. PayCalc.
Q. P-A-Y-C-A-L-C?
A. I'm not sure about that.
Q. How about PayConf, C-O-N-F?
A. I'm sorry. I don't remember that one, either.
Q. So you'll know what we did is you produced to us your -- AutoZone produced to us about 700 e-mails. And so what I'm trying to do is take some of the information that we drew from those e-mails and ask you questions in a simpler form than showing you, you know, 50 different e-mails. Okay?
A. Okay.
Q. So we believe that that's what these e-mails are saying, and I'm just trying to confirm from your knowledge whether I'm accurate or not.
A. Okay.
Q. PayConf, you're not familiar with, either?
A. No, I'm not.
Q. How about PayCalen, C-A-L-E-N?
A. No. I don't remember that at all.
Q. How about PQueue.c, P-Q-U-E-U-E, dot C.
A. I don't know what that is.
Q. How about PProdemon.x? I'll spell it. It's P-P-R-O-D-D-E-M-O-N dot X.
A. There should only be one D in that. That one,

I'm familiar with. That one uses -- it uses Semaphores. It uses shared memory. I don't recall whether it uses message queues or not. I don't think it does.
Q. How about Parts Pro?
A. Yes. That one uses shared memory and Semaphores. It works with PProdemon. And I don't know whether it uses message queues or not.
Q. What do those programs do?
A. Parts pro is our parts lookup system. Pprodemon was originally written to enhance the functionality of Parts Pro to make it faster. Parts Pro demon's job was to load some of our smaller databases into shared memory so that data could be accessed quickly. It was also originally supposed to monitor the contents of that memory, looking for errors that might develop.
As I've said, we have some power problems. We have a lot of computers in a lot of harsh places. And anything can happen. I didn't want to have corrupt memory impacting store operations. So this demon loaded stuff into shared memory and also computed check zones, so that if an error was ever found, it would tell parts lookup to go ahead and use data files because we're not sure about the contents of shared memory anymore.
Q. So this is something that worked in

conjunction with your parts lookup software?
A. That's right.
Q. And the parts lookup software was what you had referred to earlier that you believe may have been shared with customers? Is that correct?
A. Parts lookup system is or was shared with Tire Kingdom to the extent that I think they made a connection into a server in AutoZone's corporate headquarters. I don't know whether we ever put a machine containing our parts lookup in a retail customer site.
Q. Okay. POS?
A. POS. Yes.
Q. That would be one that had some Semaphores, message queues, and shared memory?
A. I believe it did use Semaphores. It may have used -- I remember discussions about shared memory, but I don't know whether that code was ever activated or used. Message queues, I don't remember whether it uses message queues.
Q. Okay. So you don't know if that code is currently being used by AutoZone?
A. The POS code is definitely being used. The code that uses shared memory and message queues and Semaphores, I'm not sure whether it currently or really

ever used it in production. I know we at some point wrote code to use shared memory to again make it faster. But I don't know if it was ever used in production.
Q. How about warranty? Is that a program?
A. Warranty is a collection of programs that allow us to keep track of our -- allows us to track of our customers purchases for warranty purchases.
Q. Does it -- I'm sorry.
A. That program does use, or set of programs, uses message queues, and I don't know about Semaphores or shared memory.
Q. Okay. How about MESGDRVR dot C?
A. I'm not certain about that program. I believe that it is part of our commercial software. I'm not -- I don't know what it does or how it does it.
Q. The programs that you and I have discussed that you do recall, are all of those in use in AutoZone at this time?
A. I think all the programs that we've talked about are still in use at AutoZone.
Q. Are there any other major programs that AutoZone uses on a daily basis that as you sit here today, you can think of that use Semaphores, message queues, and shared memory?
A. I'm sure we have other programs that use some

of those capabilities. One I can think of that uses message queues is the software that we used to communicate data to and from our stores. We talked about the FEP machines.

Q. Uh-huh (affirmative response).

A. The software that runs on those FEP machines uses a message queue.

Q. Does it have a name, this software, or you just don't remember?

A. It's called the FEP software. Its name is AZ FEP.

Q. Okay. Anything else that as you sit here today that you can think of?

A. I would not be surprised to find that other programs do. But I'm not -- I can't think of any others right now.

Q. So to your knowledge with respect to any of the implementations that we've talked about, the various versions of Linux that we talked about and the various binaries which were ported to Linux, to your knowledge, did anybody at AutoZone consult copyright counsel or consider the intellectual property implications of what it was doing when it was making these implementations?

MR. STEWART: Object that that's outside the notice, but you can answer.

THE WITNESS: Can you read that back, please?

MR. STONE: Thank you for not asking me to ask it again.

(WHEREUPON, THE LAST QUESTION WAS READ BACK BY THE COURT REPORT.)

MR. STEWART: And I'm also going to object to the form. You can answer.

A. I don't know that we looked into those issues because we were porting the software. That's it. I don't know that we looked into those issues.

Q. (BY MR. STONE) How about when you decided to use these various versions of Linux? Did you -- to your knowledge, did you look into the issues of whether anything in Linux that you were going to use may have IP implications?

MR. STEWART: Same objections.

A. No.

Q. (BY MR. STONE) To your knowledge, does AutoZone have internally any code of ethics about intellectual property or, you know, checking to make sure that before it carries out any operation pertaining to software or computers, that it doesn't violate intellectual property rights?

MR. STEWART: Same objections.

A. I'm sure we have a code of ethics that prohibits the unlicensed copying of licensed software.

Q. (BY MR. STONE) Other than that, any other code of ethics that talks about steps that need to be taken to insure that you don't copy unlicensed software or invade intellectual property rights?

MR. STEWART: Same objections. We're getting pretty far off the ranch here.

MR. STONE: This is my last question on this issue.

A. Aside from prohibitions against copying, I'm not familiar with specific processes, policies or procedures.

MR. STONE: Okay. Let's take another five-minute break. I think I have a couple more pieces of areas I want to cover, but I want to do it efficiently. I've got to break and get the pieces together.

MR. STEWART: Fine with me.

MR. STONE: Let's take a break.

VIDEO SPECIALIST: Off the record 14:35 .

(SHORT BREAK)

VIDEO SPECIALIST: Back on the record at 1459.

Q. (BY MR. STONE) Mr. Celmer, when you --

when you made your -- AutoZone made its migration to Linux, did it employ -- other than the Red Hat people, did it employ any outside consultants or contractors in making that transition?

A. I don't think so. No.

Q. So you didn't hire any consultants to come in and show you how to do it, for example?

A. No.

Q. Do you know if you hired any programmers on a contracting basis, or was it all in-house programmers that did it?

A. We have a relationship with a couple of contract programming companies, but I don't think that any of those guys were involved with the port because they're usually here for specific assignments. And the porting of the software for a long time was not really an organized effort. And once it became an organized effort, we were sort of in a hurry to do it. So I don't think we brought -- tried to bring the contractors in on the port.

Q. So based on your prior testimony, am I -- am I correct that you began doing this porting process informally in between other projects that the IT department was doing?

A. Right.

Q. And Mr. Greer was the person that was primarily responsible for sort of supervising that informal process? Is that correct?
A. That's right.
Q. And then it became more and more involved as time went on?
A. That's right.
Q. But it still was never that you picked a certain group of people and said you are the Linux development team and you will do nothing but work on porting applications from OpenServer to Linux. Is that correct?
A. Toward the end of the porting effort, it became a sure enough project where people were dedicated to it on a full-time basis. And so at that point, there was a team, but they were all AutoZoners. I don't think there were any contractors.
Q. Was this before Mr. Greer left or after he left?
A. Before he left.
Q. When you cre -- when AutoZone created the binaries -- strike that. When AutoZone ported binaries it was using on OpenServer to Linux, did it use the standard errors -- error messages that were provided in Unix?

A. I don't think I understand what standard error messages you're talking about.
Q. Are you familiar with a ERRONO dot H header fire -- header file?
A. Yes.
Q. I'm glad you are. That's what I'm talking about. Did you use the error messages that are referenced in that file when you ported binaries?
A. Yes. I'm sure that we have some programs that refer to the standard names for those error numbers. And those names were available under Linux just as they were under OpenServer for the most part. There are a couple of differences.
Q. So now that you -- your store machines are running on Linux, they're still using many of those same error messages? Is that right?
MR. STEWART: Object to form. You can answer.
A. We're still using the same -- I'd rather say names than messages. Because I believe there's a separate header file that provides, you know, multi-word messages. And ERNO dot H is a simple mapping of a number to a standardized name, one word, just a name. So changing messages to name, then the answer is, yes, we're still using those names.

Q. (BY MR. STONE) So like error -- error no dot H would have in it like the number one, two, three, or something that would refer to an error message?
A. That's right.
Q. So like those annoying messages you get on your computer when it says error number 212 or something, and you have no idea what it means, that's the kind of thing we're talking about?
A. Similar.
MR. STEWART: Except Bob does know what it means.
Q. (BY MR. STONE) When you ported binaries -- when AutoZone ported binaries from OpenServer to Linux, did it use the signal numbers that were defined in the signal dot H header file?
A. Yes.
Q. And the 3500 stores that are using the Linux operating system, do they still use those signal numbers?
A. Yes.
MR. STONE: Would you mark this, please.
(WHEREUPON, THE ABOVE-MENTIONED DOCUMENT WAS MARKED AS EXHIBIT NO. 38,

AND IS HERETO ATTACHED.)
Q. (BY MR. STONE) I'm showing you what's been marked Exhibit 38 for identification.
MR. STEWART: Do you have an extra copy?
MR. STONE: No. I didn't have a copy for myself.
MR. STEWART: Let me take a look at it first.
Q. (BY MR. STONE) I'll represent to you it's a printout of the posting that Mr. Greer posted on Groklaw that I referred to earlier. Could you -- could you review it, please? I'm going to ask you some questions about it.
(BRIEF PAUSE)
A. Okay.
Q. (BY MR. STONE) Have you had an opportunity to read it?
A. Yes.
Q. Is that the first time you've ever read that?
MR. STEWART: And I'm going to object, by the way, to this document to any questioning about the document on the grounds that it's outside the deposition notice. But I'll allow the witness to answer.
A. I think I've seen parts of this, in particular

the first sentence. I seem to remember it.
Q. (BY MR. STONE) Could I have it back for a second?
A. Sure.
Q. The sentence, I don't know whether to be pleased or angry at SCO's assertion that IBM must have assisted AutoZone's transition to Linux, et cetera?
A. Right.
Q. Was -- was Mr. Greer a senior technical advisor at AutoZone?
A. Senior technology advisor, yes.
Q. That was your title, as well?
A. And is currently. Yes.
Q. Did you note in here where Mr. Greer states that as many as 35 AutoZone developers were performing porting work for him?
A. I did note that.
Q. Does that refresh your recollection? I believe you told me earlier you thought there were no more than 12 to 15 developers working on this project at any one time.
A. Yeah. That number surprised me because that's a lot of folks.
Q. So as you sit here today, do you know whether your testimony is more accurate or his testimony as to

how many developers were working on this project?
A. Well, there may be some middle ground, as well. It may be that there were a total of 35 people on, you know, a particular day or set of days. But it may not necessarily be the case that that's the only thing they were working on.
So I was talking about full-time dedicated people. It may be that he's talking about a short-term requirement for everyone to pitch in and do porting at the same time as testing or some other activity. But this was -- the porting of the software was his responsibility. And so his recollection is almost certainly more accurate than mine.
Q. Mr. Greer stated that the developers were also responsible for testing their individual applications under both SCO and Linux. Is that what you were referring to earlier when you were talking about having the machine on your desk and referring to a development machine such as Vision and Wrangler and running the binaries on the two machines?
A. Yes. That's exactly it.
Q. And comparing the output?
A. Right.
Q. And then he refers to performing bills of the entire AutoZone store on his desktop. Is that the

Wrangler machine, or is that some other machine, if you -- if you know?
A. I really don't know. Jim was, around that time, our primary investigator of Linux. He had a keen interest in Linux. And in fact, his desktop was a Linux machine. His desktop might have been Wrangler for a period of time while he was doing this on his own. And when other people began to join his efforts, he had to make that machine available. So he could have made it available from his desktop, but ultimately he would have had to separate that machine from the developers, you know, physically and replace his desktop with something else.
Q. His statement that no SCO libraries were involved in the porting activity you now know to be incorrect. Correct?
A. Well, if a program was ported, then there weren't any SCO libraries involved. But our problem is with those files that were copied but not ported.
Q. So the point is that there were a number of files that were copied and ended up on your 3500 store machines that were not recompiled under Linux and were not ported in accordance with the structure that you had set up. Is that correct?
A. That's right. The problem with those programs

is they weren't ported, or if they were ported, Linux binaries weren't distributed.
Q. Mr. Greer also refers to rules that he had established for these developers. Would it be fair to say that to your knowledge, he didn't establish any rules that pertained to protections of SCO proprietary code?
A. I believe his overwriting rule was port everything, which had that been executed properly, would have protected IB -- excuse me -- SCO's library code from being copied.
Q. But you've now learned that it wasn't done correctly?
A. That's right.
Q. When you were considering -- considering the migration of AutoZone to Linux, were you -- did you rely on any certifications from IBM that related to Linux?
A. I don't think so.
Q. Did you have knowledge that there were certifications that IBM had generated pertaining to Linux?
A. No. Because I'm not even familiar now with what certification IBM might be able to provide.
Q. Were you aware of IBM announcing in the 1999/2000 time frame that it was embracing Linux as

planning to spend billions of dollars on advancing Linux?

A. I'm unsure of the time frame, but I do remember IBM promoting Linux as it promoted Java.

Q. And did that play any role in your decision to migrate from OpenServer to Linux?

A. No.

MR. STEWART: Object. You can go ahead and answer.

A. No. I'm not really concerned with what IBM thinks people should do.

MR. STONE: Okay. Take a break. I think I'm done. I may have one or two more questions. But I think we're pretty close to done at this point.

VIDEO SPECIALIST: Off the record, 1515.

(SHORT BREAK)

VIDEO SPECIALIST: Back on the record at 1521.

Q. (BY MR. STONE) Mr. Celmer, do you know what an INIT tab is?

A. Yes.

Q. Can you tell me what it is, please?

A. There is a program called INIT that is started at the end of the process of booting up the Unix machine, and that program requires some information so

that it knows what it's supposed to be doing. And so there's a table of information that is called a INIT tab, which, I suppose, is for an INIT table. And that table of information usually tells INIT which programs to start, in what order, at what point, and at which point of the boot process it should start them, and it tells them whether they should restart them if they die and some other information.

So it's -- the file is used by this program to make sure that when a system comes up, all of the services that are necessary at that particular run level are running or at least that they try to start.

Q. Have you heard of something called an INIT script?

A. INIT script. I believe I have heard of INIT script, but I don't know what the heck it is.

Q. So you don't know if it has something to do with INIT tab or not?

A. I'm sorry. I don't know what it is.

Q. Okay. Do you know if AutoZone has in any of its binaries that are in its 3500 stores on its Linux systems, INIT tabs?

A. The process of booting Linux is similar to that of booting Unix. And there is a program that performs a function that's similar to INIT. I believe

in our version of Linux, that program is called EX-INIT. But I don't remember whether we're using an INIT tab file or whether there's some kind of EX-INIT dot conf file, which is something that Linux does often. They put configuration in the file and put a dot conf at the end.

So I'm sorry. I don't really remember whether we have an INIT tab file in our stores or not.

Q. Did you have INIT tab files in your stores when you were using OpenServer?

A. I believe we did. Yes.

Q. When you ported your applications from OpenServer to Linux, do you know what you did about those INIT tabs? Is that something you would have taken out or replaced, or how would you have dealt with it?

A. Well, I don't recall directly what was done. But among the things that are in the INIT tab file are the instructions for how to allow users to log into terminals, and we certainly would have needed that functionality. So we would have instructed this INIT or EX-INIT program under Linux to start the program that communicates with terminals and allows the user to log in.

There are a number of other things that INIT does. So we would have reviewed what INIT was instructed to do

via the INIT tab file under Unix, and we would have tried to replicate that functionality for the corresponding program under Linux.

Q. Would you have done that in the way that you've described before, using the two different development machines, the OpenServer machine and the Linux machine, to try to replicate the same function?

A. No. I think because Jim ported just a couple of programs to start with, that he would not have had to bring up a full system sufficient to run all of our software. All he wanted to do was recompile the source of the libraries, which took a while, as I remember, and then the source for, you know, some set of programs. And just in order to run those programs, he would not have had to modify the INIT file because he was working at his desk, which would have been the system console. So I don't think he would have had to modify INIT until well into the project.

Q. So it would have been during when you were getting close to rolling out the project that you would have done that? Is that correct?

A. It wouldn't have been necessary until then. I really don't know at what point of the project he would have done that. Because you know, first, we have to make sure that programs run and this is a viable thing

to invest time in. But at some point along the way, even if all your programs run, if you can't get the system up, it's not worth porting the programs.

So he would have had to evaluate at what point to investigate that critical piece, and I really don't know when he did it.

Q. But as -- as you sit here today, you don't know when he investigated that critical piece or when he actually performed the porting activity. Is that right?

A. That's right. I don't.

MR. STONE: All right. I don't have any further questions.

MR. STEWART: No questions.

MR. STONE: Thank you for coming.

VIDEO SPECIALIST: This concludes the videotaped deposition of Mr. Bob Celmer, consisting of four tapes. The time is approximately 1528. We are now off the record.

CERTIFICATE

STATE OF TENNESSEE:
COUNTY OF SHELBY:

I, KORIAN NEAL, Registered Professional Reporter and Notary Public for the State of Tennessee at Large, do hereby certify that I reported in machine shorthand the above-captioned proceedings.

I HEREBY CERTIFY that the foregoing pages contain a full, true and correct transcript of my said Stenotype notes then and there taken.

I FURTHER CERTIFY that I am not an attorney or counsel of any of the parties, nor a relative or employee of any of the parties, nor am I a relative or employee of any attorney or counsel connected with the action, nor am I financially interested in the action.

I FURTHER CERTIFY that in order for this document to be authentic and genuine, it must bear my original signature and my embossed notarial seal and that any reproduction in whole or in part of this document is not allowed or condoned and that such reproductions should be deemed a forgery.

THEREFORE, witness my hand and my official seal in the State of Tennessee on May 16, 2005.

_____
KORIAN NEAL, RPR, CCR
Notary Public at Large

My Commission Expires:
June 7, 2006