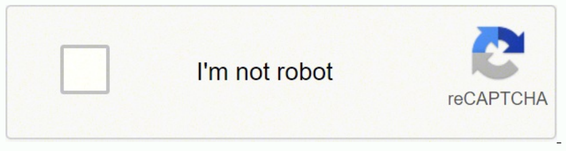

Continue

14082339.044944 16777165.876923 170674376238 12546999.288889 82886890906 39180558432 26578018.123077 86695753.611111 161038625996 100271724990 3877453846 11490085.159574 68048645388 92814395781 27345763.897059 6961129164 699155.15384615 43643732880 10134798995 53765454005 30802392520 16771778774 58210443393 9062955.3444444 8878533.7777778 820686.71276596 39186287.311111 20605876.75 365602698 52690294656

A practical guide to enterprise architecture

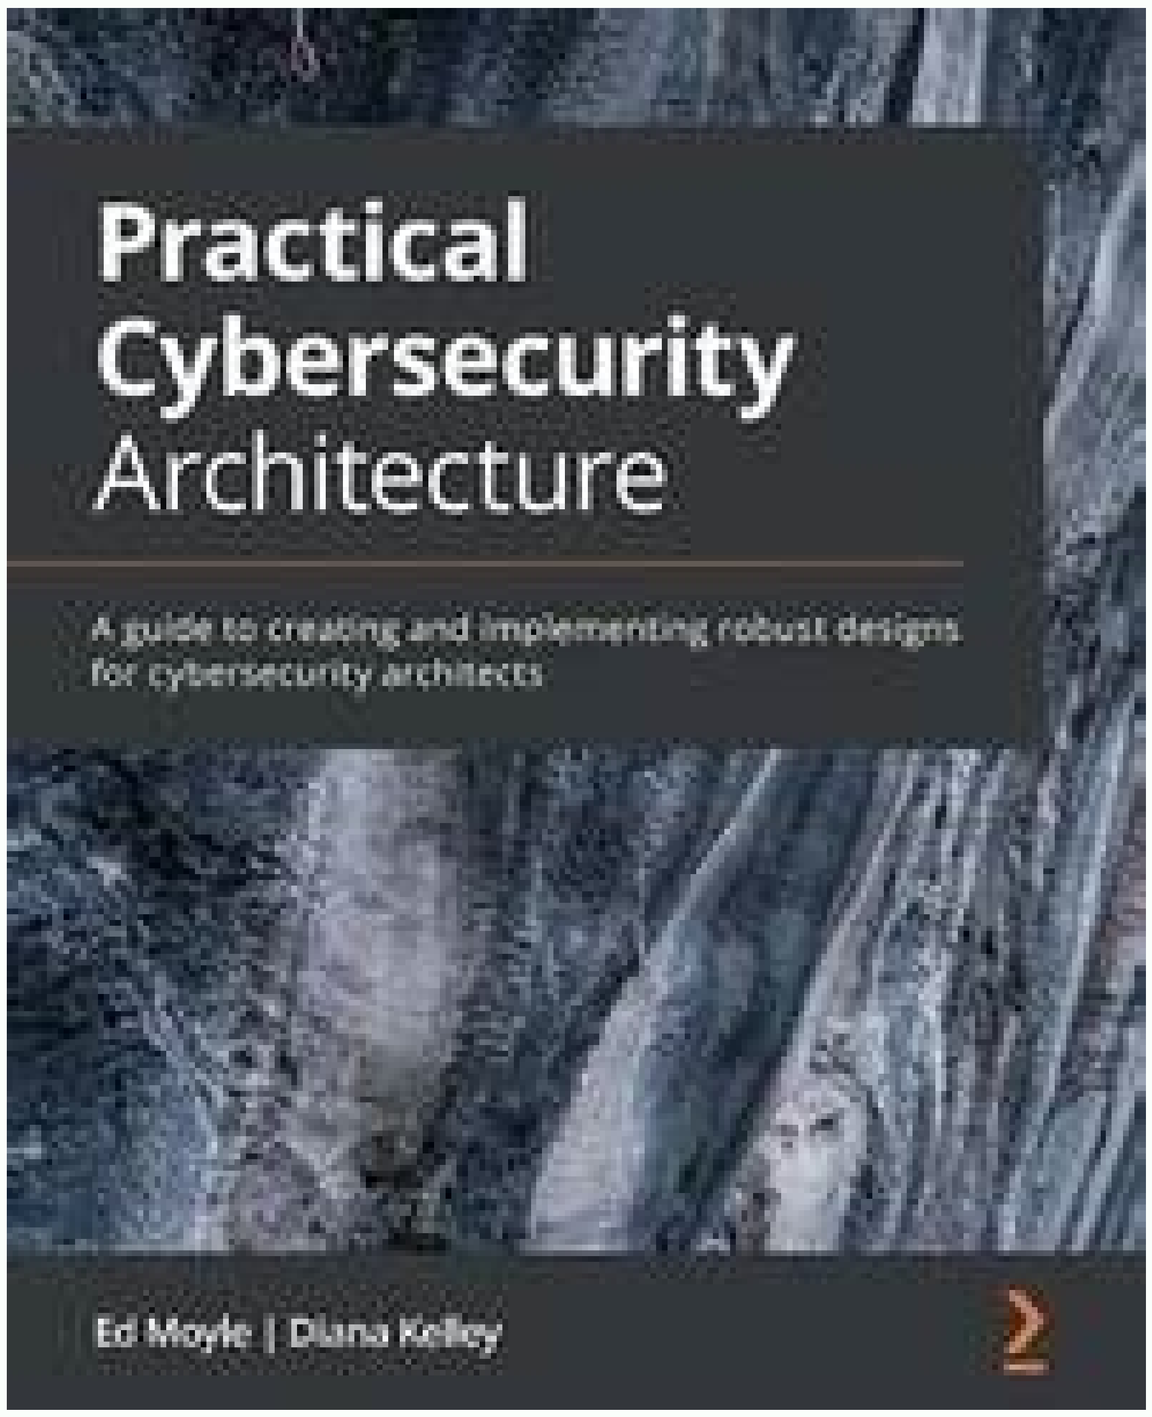

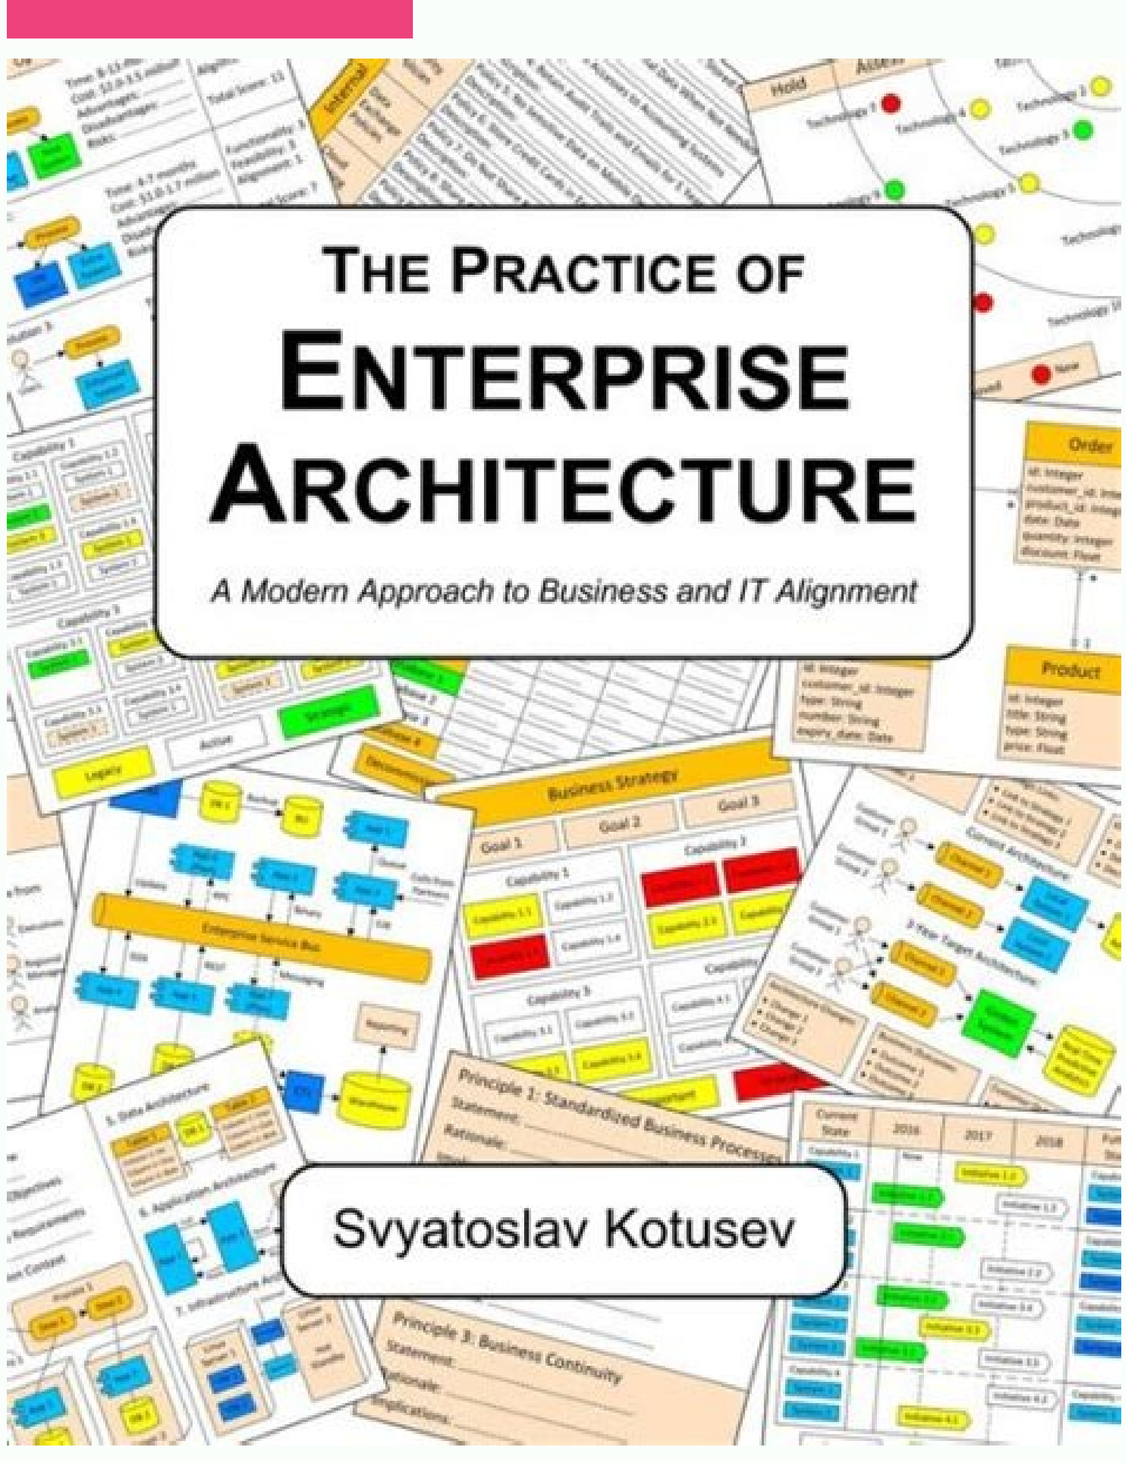

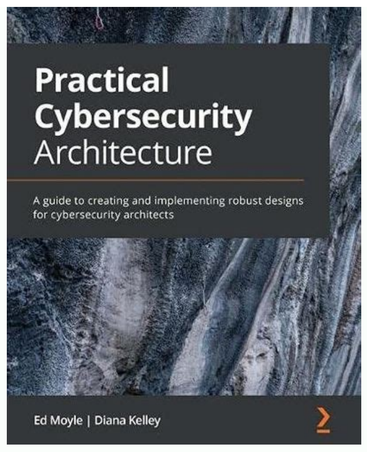

A practical guide to enterprise architecture. A practical guide to federal enterprise architecture. A practical guide to enterprise architecture pdf.

Although not every attribute is basic, and this section looks at the more advanced concepts. Therefore, the user's password and the registry are needed to decrypt any encrypted files that are encountered during an investigation. Entries 17 to 23 are sometimes used as overflow when the reserved entries are not enough. We read the root directory contents from block 1096 and process the contents as a list of directory entries. The directory at inode 32577 is the most promising. We will now examine each of these data structures in more detail. To determine the allocation status of an inode, we use the inode bitmap for the group, which is located in the group descriptor. Three scenarios of NTFS indexes, including A) a small index of three entries, B) a larger index with two nodes and 15 entries, and C) a three-node tree with 25 entries. This is shown at the bottom of Figure 11.15. The restriction is typically placed at 25% or 33% of the total number of blocks. Standard Attribute Types So far we have been speaking in general terms about attribute types. With UFS, there is a lot of non-essential data in this category, and detecting hidden data is further complicated because the essential and non-essential data are intertwined. When access is revoked from a user, her key is removed from the list. We can see this in Figure 11.19 where we have an $INDEX_ROOT attribute with three index entries and a non-resident $INDEX_ALLOCATION attribute that has allocated cluster 713, and it uses three index records. For example, the base might increase by 32 fragments for every group and then start at an offset of 0 after 16 groups. The file system tools in TSK are based on the tools from The Coroner's Toolkit (TCT) ( ), which is by Dan Farmer and Wietse Venema. An MFT entry can have multiple attributes of the same type. For example, consider MFT entry 313 with a sequence number of 1. For example, one attribute is used to store the file's name, and another is used to store the file's content. To decrypt a $DATA attribute, the $LOGGED_UTILITY_STREAM attribute is processed and the user's DDF entry is located. . The dstat tool in TSK shows us the allocation status of a UFS fragment and gives us the cylinder group that it is a part of. This section covers where the data are located in UFS and how they can be analyzed. It shows one deleted file, the file with the "*" before it. The last example was too simple, so I will present the more challenging file shown in Figure 11.10. The rest of the times are left as an exercise, if you are really bored. The extended attributes also might have evidence or hidden data and should be examined. Those changes are not included in this description, but 2.00 could be available by the time you read this. Also notice that block 384 is being used as an indirect block pointer. VDOC.PUB Download Embed This document was uploaded by our user. . NTFS uses VCN-to-LCN mappings to describe the non-resident attribute runs. In both of these scenarios, we can still determine that the middle file was the last to be deleted. The final unit compresses to 16 clusters, so it is not compressed and a run of 16 clusters is created. The next chapter shows example data structures that contain symbolic links. 4th ed. MFT Entry Contents The size of each MFT entry is defined in the boot sector, but all versions from Microsoft have used a size of 1,024 bytes. To find the inconsistencies, an OS runs a program that scans the file system and looks for missing pointers and other signs of corruption. "Overview of FAT, HPFS, and NTFS File Systems." Knowledge Base Article 100108, 2003. The starting location of each of these data structures is given as a byte offset that is relative to the start of the group descriptor, and the size of the data structure must typically be calculated. fls will list the file names in a given directory. In addition to the group descriptor, each group also contains an inode table and a backup copy of the superblock. This was used in several of the scenarios in Part 3, "File System Analysis," of the book. Our analysis tool might show that the bbb.txt file was deleted, but it really wasn't. Analysis Considerations Deleted file names are easy to locate in ExtX, and the Ext3 inode number is not cleared by Linux, so you might also be able to obtain temporal information about when the file name was deleted. File System Journals As any computer user knows, it is not uncommon for a computer to halt and crash. If there are multiple $DATA attributes in the MFT entry, they are all encrypted with the same FEK. UFS will wipe only the fragments allocated, so parts of a block will still exist after some of it has been reallocated. To do so, we should examine the first 24 bytes of each block to determine if the '.' and '..' entries exist. If an incident occurred very recently, you might be able to draw conclusions about where files were deleted from and which file was last allocated. Throughout this book, we have differentiated between the different types of addresses. Attribute Content The content of the attribute can have any format and any size. Several security tools can perform a brute force attack against a user's login password, and this can be used to decrypt the data. It is trivial for an executable to obfuscate the names of the files that it opens. On the outside of the box is basic information, such as your name and address. There are several data structures that keep lists of available fragment sizes and clusters that are ignored for simplicity in this example. Figure 11.1 shows the basic layout of an MFT entry where there is some header information and three attributes. The file system tools are further organized into the data categories that we discussed in Chapter 8, "File System Analysis." Each tool name has two parts, where the first part identifies its group and the second part identifies its function. The other contains the file systems and has a Volume Table of Contents (VTOC) data structure in sector 0. In Figure 11.5 we see the example MFT entry that we saw previously, but now its third attribute is too large to fit in the MFT, and it has allocated cluster 829. A journal shows which file system events recently occurred, and this could help with event reconstruction of a recent incident. The time value is also a 16-bit value and also has three parts. Therefore, before we determine the allocation status of a block, we will need to determine in which group it is located. Microsoft reserves the first 16 MFT entries for file system metadata files.[2] The reserved entries that are not used are in an allocated state and have only basic and generic information. In an abstract sense, there is no essential data in this attribute, but the application-level features of the file system require it to be there. 176 $BITMAP A bitmap for the $MFT file and for indexes. If we encrypt a file with symmetric encryption and want multiple people to access it, we need to either encrypt it with a key that everyone knows or make a copy of the file for each user and encrypt each with a key that is unique to that user. If the record length of an unallocated directory entry is the actual size that it needs (based on the length of its name), the next directory entry was deleted after it was. Overview UFS uses fragments and blocks to store file and directory content. Boston: Addison-Wesley, 2004. The swap space or page file might also provide copies of unencrypted data. An i386 Solaris system is similar where it will have two DOS partitions. Group 2 starts in block 65,552, so its inode table starts in block 65,608. 64 $VOLUME_VERSION Volume information. "How to Fix Cross-linked Files." Microsoft Knowledge Base Article83140, May 10, 2003b. The second unit is all zeros, so a sparse run of 16 clusters is made for it, and no clusters are allocated. Deleted files have the block pointers, size, and mode cleared from the inode, but UFS preserves the state of the indirect block pointers. Any fragment that is not allocated is extracted for analysis. We read the inode table from block 56 and process the third entry (the first entry is inode 0). The consecutive fragments cannot cross a block boundary, and the bookkeeping information in the file system provides a list of where fragments of a given length can be found. Boston: Addison Wesley, 2005. For example, if we want to find the value 6, we compare it to the root value, which is 7. The address of the block and the starting fragment are added to the inode. The base address is a variable number of blocks from the start of the group, and the bitmaps are a static distance from the base. Sometimes a backup is the only available data, and the investigator needs to make the most of it. The last relevant data structure is located in the cylinder group summary area. Using this information, our theory is that snifferlog-1.dat file was created after the only_live_twice.mp3 file, and then lic_to_kill.mp3 was created. Chris rated it really liked it Aug 30, 2014 Charlie rated it it was ok May 25, 2013 Hjavadi rated it liked it Mar 23, 2015 Rajesh rated it it was amazing Feb 16, 2016 Jan 08, 2009 James rated it it was amazing · (Review from the author) Michael rated it liked it Apr 18, 2014 Mike van Alst rated it did not like it May 06, 2018 Stefan rated it liked it Sep 13, 2014 Chris rated it liked it Sep 26, 2020 Worthy rated it it was ok Dec 01, 2013 John rated it liked it Oct 14, 2019 Avcastoldi rated it did not like it Jul 26, 2013 Dean marked it as to-read Oct 29, 2012 Excelsior marked it as to-read Mar 29, 2013 Simon marked it as to-read Oct 13, 2013 Kaliprasad marked it as to-read Jan 30, 2014 Sripathys marked it as to-read Apr 25, 2014 Loading PreviewSorry, preview is currently unavailable. See also the Bibliography section of Chapter 9. The UFS superblock is located somewhere in the start of the file system. "Inside Win2K NTFS." Part 1, Windows and .Net Magazine Network, November 2000. 96 $VOLUME_NAME Volume name. Attribute Headers The attribute header identifies the type of attribute, its size, and its name. Soft links are also a second name for a file or directory, but they can span different file systems. This is easy because all we need to do is divide the fragment address by the number of fragments per cylinder group. For example, the next example lists the allocation status of each data unit in an NTFS image: # dls f ntfs e l ntfs-10.dd addr|alloc 0|a 1|a [REMOVED] 13423|a 13424|f The 'a' after each address signals that the data unit is allocated, and an 'f' signals that it is unallocated. Figure 11.11. He co-authored Java Web Services Architecture. There could be an allocated metadata structure with allocated data units, but no pointers between them and no file name pointing to the metadata structure. softlink.txt has its own inode that contains the path of the file. This can be seen in Figure 10.3(A). The time updating for OpenBSD 3, FreeBSD 5, and Sun Solaris 9 systems are the same as reported in ExtX for Fedora Core 2. They are limited to two blocks in size, but test if your analysis tools will show you this content and if they include the content in a keyword search. Only the 1-3-2 and 2-3-1 sequences have the same final state. The type of partition table can be specified on the command line using the -t argument and the type, which are given in this paragraph in parentheses. Inodes Inodes in UFS have the same basic concept as we saw with ExtX. The valid range of this value is 0 to 29, which allows a second range of 0 to 58 in two-second intervals. Brouwer, Andries. To make the scanning program's job easier, some file systems implement a journal. After we advance to the boundary, we apply the directory entry data structure and perform sanity checks to determine if the data could have been for a directory entry. It compares that with the actual record length. 2 4,084 The final '..' entry has a record length of 4,084 bytes because it needs to point to the end of the block, but it needs only 12 bytes. If we acquired at the volume level, the hidden data would be lost. The BSD and Solaris systems that I tested will all wipe the unused sectors in a fragment. A scenario where a backup would be critical is in a corporate environment where a server is not responding because its disks were wiped with 0s and then rebooted. Similar work has likely occurred with the commercial file systems, but the work has not been published, and my basic testing might not have shown the full extent of their algorithms. Microsoft MSDN Library. If we do not think that the IDS was compromised, the only evidence on the system is at the file level, and we can simply copy the necessary logs and take the appropriate preservation steps. Instead of saving names and addresses in a file, they would be saved to a special section of the volume. To locate the data, we need its fragment or block address. The basic directory entry structure has the fields given in Table 10.5. Table 10.5. Data structure for a basic FAT directory entry. A sparse attribute is one where clusters that contain all zeros are not written to disk. This section gives an overview to each of the tools in TSK. 0000048: 0d00 0000 0e00 0000 0f00 0000 1000 0000 ..... A directory entry is a simple data structure that contains the file name and the inode address where the file's metadata can be found. To determine which group an inode is in, its address is divided by the number of inodes per group, which can be found in the superblock. When a file is extended and it already has fragments, the OS first tries to extend the existing fragments. Content Category The content category of data includes the file and directory content. Even if the dir1 directory has its own files in FS1, they will not be shown when FS2 is mounted on it. We used this command in the NTFS chapters because it stores all data in files. For investigators, this means that you need to know where file systems were mounted. The remaining bytes store attributes, which are small data structures that have a very specific purpose. "Defining Digital Forensic Examination and Analysis Tools Using Abstraction Layers." International Journal of Digital Evidence, Winter 2003a. . Inodes 0 and 1 are reserved, but not used for anything. The group descriptor will contain the fragment bitmap. Extended attributes are stored in normal data blocks, and the block addresses are given in the inode. The data structures for the node descriptors and an actual directory are given in Chapter 15. In other words, a directory can store both file content and a list of its files and subdirectories. 2 12 File1.dat 8 4,072 There are actually two versions of directory entry structures. An OS can choose any algorithm, but the typical policy is to allocate an inode for a directory in a new cylinder group that has a less than average number of directories and a greater than average number of available blocks. The first block and fragment start with the first sector of the file system. Other OSes could choose a different strategy. . The boot code is located in sector 0 and then in sectors 2 to 15. The rest of the first block contains node descriptors, which contain a hash value and a block address. Each block contains a list of variable length data structures that have the fields shown in Table 17.9. Table 17.9. Data structure for the UFS2 extended attribute entry. The unused 32-bit fields are ignored and not reused. Inside Windows 2000. The example shown is a binary tree because there are a maximum of two children per node. The latter attribute exists for every file and directory because it contains the data needed to enforce data security and quotas. We will be using this example image later in this chapter, and it is used for the manual analysis in the next chapter. 0000048: 2000 0000 0104 0464 6174 6500 0000 0000 .....date..... Data Acquisition Layers The general theory of non-volatile data acquisition is to save every byte that we think may contain evidence. George, Esther. For example, if our NTFS file system had 4,096-byte clusters and we found evidence in the 123rd cluster in the unallocated data file, we would supply 123 with the -u flag: # dcalc f ntfs u 123 ntfs-10.dd 15945 We also can determine the allocation status of a specific data unit by using the dstat tool. Every directory has an $INDEX_ROOT attribute that contains information about the files and subdirectories that are located in it. A data recovery field is created for each method of data recovery, and it contains the FEK encrypted with a data recovery public key that is used when an administrator, or other authorized user, needs access to the data. For example, A is a parent node to B and C, which are children of A. Although only privileged users can read the file's system name space attributes. The block pointers are 32-bit values in UFS1 and 64-bit values in UFS2. Figure 11.8. A 12-cluster file that is stored in A) normal layout and B) sparse layout with a sparse run of three clusters. Table 17.10. We read the inode table from block 65,608 and process entry 37, which is the relative location of inode 16,549. The superblock defines how many blocks should remain free at any given time. Prior to version 3.0 of NTFS (which came with Windows 2000), only the $FILE_NAME attribute was in an index, but now there are several other uses of indexes and they contain different attributes. The difference between the UFS1 and UFS2 superblocks is that the UFS2 version includes 64-bit versions of the size and date fields, which were added to the end of the data structure. The second one is because every cylinder group has a backup copy located at an offset of 16 fragments. In Linux, the directory entry structure will remain in the unallocated state until a new file is created whose name is the same length or smaller. There are many types of attributes, and each has its own internal structure. Page 16 Brenner, Susan, Brian Carrier, and Jef Henninger. Figure 11.13. When stored as a sparse attribute, three runs are created and only nine clusters are allocated, which can be seen in Figure 11.8(B). Refer to the man pages or the website for more details. A disk for a Sparc Solaris system has a VTOC in sector 0 of the disk and the boot code in sectors 1 to 15. We first need to process the group descriptor to find the offset of the fragment bitmap, which is located at byte offset 1,200. Microsoft does not delete MFT entries after they have been created. These data are not essential to the file system, and they typically exist as special file system data instead of inside a normal file because it is more efficient. "Windows Server 2003 Technical Reference." Storage Technologies Collection Section, 2004. Lastly, there is a hash database tool named hfind, that allows you to quickly lookup a MD5 or SHA-1 hash value from the NIST NSRL or one that you made using md5sum. . Full blocks are allocated when the file is extended. 112 $VOLUME_INFORMATION File system version and other flags. The entire file system is considered a data area, and any sector can be allocated to a file. Entry File Name Description 0 $MFT The entry for the MFT itself. Figure 11.12. Carrier, Brian, and Eugene H. Notice that in Part B both inodes have a link count of one. Examples of this were given in each of the previous file system chapters. If we think of our boxes analogy, there is always the same basic information on the outside of each small box, but the shape of each box may be different. To locate a specific inode, we need to first identify its group, and we can do that by dividing the inode address by the number of inodes per group. Tree after deleting the 'zzz.txt' file and the 'fff.txt' file. Sparse Attributes NTFS can reduce the space needed by a file by saving some of the non-resident $DATA attribute values as sparse. A directory entry has a dynamic length because the file name can be anywhere from 1 to 255 characters long. In this section, I will give an overview of the documented BSD allocation strategies but would recommend the Design and Implementation of the 4.4 BSD Operating System or Design and Implementation of the FreeBSD Operating System books for anyone who wants more details. A first-available strategy is used when allocating the final fragments. In theory, an OS could create a fixed number of entries when the file system is created, but the dynamic nature of Microsoft's implementation allows it to easily make the file system larger when more space is added from volume spanning. An empty entry is used to signal the end of the list. This is similar to how an MFT entry has no internal structure and it contains several attributes that contain specific information. This creates a tree that is still balanced where all leaves are the same distance from node H. By default, ils will show only unallocated metadata entries, but all of them can be shown with -e. "FAT: General Overview of On-Disk Format." FAT32 File System Specification, Version 1.03, December 6, 2000. The file content of file1.dat is written to the allocated block and fragments. NTFS Index Attributes Now that we have described the general concept of B-trees, we need to describe how they are implemented in NTFS to create indexes. "Change Journals." 2004. If it is an existing file, restrictions might be placed on how many blocks any file can have in a single group, so a new group might need to be selected. Therefore, the new entry will be added 12 bytes after the start of the '..' entry, and the record length of the new entry will point to the end of the block. . They are free to allocate in any method they want. The point of walking through the process of adding and deleting values from the tree was to show how complex the process could be. There are currently 13 tools in the file system layer tools, and they are organized into five categories. Computers do the same thing and copy data from the suspect systems in chunks of data, ranging from 512 bytes to many thousands of bytes. There is unused space between each of the directory entries, so simply advancing to the next 4-byte boundary after each entry would not find the names. Inodes that are all zeros are typically ones that have not been used and are still in an initialized state. Obviously, a real situation would use keys that are much longer than "spot" and "felix." In fact, they are typically over 1,024-bits long. In addition to the more detailed descriptions in Chapter 12, the data structures for many are given in Chapter 13. Depending on the implementation, the details of zzz.txt file may still exist in the node and could be recovered. This is one of the ways that NTFS is very different from other file systems. Know Your Enemy. For every working architect and every IT professional who wants to become one. When a file or directory is deleted, its name needs to be modified so that the OS does not print it. In either case, the OS tries to add the name in the unused area. The motivation for having two types of data units is to allow files to allocate large amounts of consecutive data while not wasting space in the last block. Microsoft says that only the $DATA attribute should be compressed, and only when it is non-resident. Using the metadata-addressing scheme of TSK, the RESUME-1.RTF file is the second entry in the root directory, which means that it has an address of 4. The header of the attribute identifies if the attribute is resident or non-resident. Notice that the types for the deleted file are still the same, which means that the inode may not have been reallocated yet. Note that the differences between the deletion routines with this and what we saw with Linux and ExtX is that Linux did not clear the mode field, but it did clear the contents of the Ext3 indirect block pointers. For example, we can view the contents of data unit 23,456 in our Ext3 image by using the following: # dcat f linux-ext3 ext3-5.dd 23456 Metadata Category The metadata category includes the data that describe a file. The starting location of the root directory in a FAT32 file system is given in the boot sector. The top of the figure gives the starting state where there are four allocated directory entries. The first consideration when allocating a block or fragment is the cylinder group. To process this file, we need to first organize all the data in the six runs and then organize the data into compression units of 16 clusters. Value Description 1 User 2 System Here we see the contents of an extended attribute block with two attributes: 0000000: 3000 0000 0107 0673 6f75 7263 6500 0000 0.....source... For example, if the system crashes at some point while the various data structures for a file are being allocated, the sequence number can determine whether a data structure contains an MFT entry address because the previous file used it or because it is part of the new file. Directory Entries An ExtX directory is just like a regular file except that it has a special type value in its inode. We will be discussing the concept of an NTFS file in this chapter. The Computer Forensic Tool Testing (CFTT) project at NIST developed requirements and test cases for disk-imaging tools. Although UFS2 added a create time value to the inode. When we examine each of these files, we notice that only_live_twice.mp3 is an executable, and the other files are network sniffer logs in the same format as snifferlog-1.dat. . The link count becomes 0, which means that the inode must be deallocated. There is also a flag that indicates if an attribute is sparse. Figure 11.7. Our example MFT entry where the type names and identifiers have been added to the attributes. To extract all unallocated fragments in the file system, we cycle through each group and examine each fragment bitmap. The superblock also might contain a volume label and the time that the file system was last mounted. For example, if the last character of a name is in byte 34, we advance to byte 36. The first 42 bytes of the data structure contain 12 fields, and the remaining 982 bytes are unstructured and can be filled with attributes. Carrier, Brian. With ExtX, we saw that the second block of each group always had a table of group descriptors. The current version has the sigfind tool, which searches for binary values. There are many variations of B-trees, and there are more rules than I will describe here because the purpose of this section is to describe their concepts, not to describe how you can create a B-tree. The final run is a compression unit and it is not compressed. The C-time of the file is updated when the touch command is used. Figure 16.3 shows an example of a UFS1 and a UFS2 file system. B-Trees An NTFS index sorts attributes into a tree, specifically a B-tree. If this is not possible, a new set of fragments or a full block is allocated and the data are moved. When this occurs, the first entry is called the base file record, or base MFT entry, and each of the subsequent entries contains the address of the base entry in one of its fixed fields. 7 $Boot Contains the boot sector and boot code for the file system. Figure 11.18. For example, $DATA is the name of the attribute type, and the attribute's name could be "fred." Some tools, including The Sleuth Kit (TSK), will assign the name "$Data" to the default $DATA attribute. See also the Bibliography section of Chapter 11. The output of mmls is sorted by the starting address of the partition, regardless of where it is located in the table. This data structure supports a name that has only 8 characters in the name and 3 characters in the extension. Therefore, a block still might contain deleted data in the fragments that have not been reallocated. The Honeynet Project. A node in an NTFS directory index tree with four index entries. Wilson, Craig. For example, the key "spot" could be used to encrypt the plaintext into ciphertext, and the key "felix" could be used to decrypt the ciphertext. 20Serial%20Numbers.pdf. We can see this in Figure 14.6(B) where the b.txt file was deleted and the pointer in a.txt was incremented to point to c.txt. d/d 53248: .. Therefore, if you have an 8,192-byte block size, each chunk will be 128 KB. We also use grep to filter out all entries that are not directories (it is using the fifth column). If the root node is smaller, we go to the child on the right. To determine the allocation status of any fragment, we must first determine the cylinder group in which it is located. . "Recovering NTFS Boot Sector on NTFS Partitions." Knowledge Base Article 153973, 2003. Table 11.1. The standard NTFS file system metadata files. Decryption process starting with encrypted content, keys, and a user password, and ending with the decrypted content. . Microsoft found that older OSes would ignore directory entries with all the bits set and would not complain about the different layout. One method is to use the file type value in the directory entry and compare it with the type in the inode. Unfortunately, node A now has four values in it. dcalc will calculate the original data unit address by using the data unit address from the unallocated data. . Listing the unallocated entries is useful to find the entries from deleted files where the file name has been reallocated. Compare the size of the file system with the size of the volume to find volume slack. /system32/ntio804.sys 35392 .a. For example, many of the OSes will restrict how many blocks a file or directory can allocate in a single group. JAMES LINN, consultant at Hartford Technology Services, co-authored XQuery Kick Start. ELIAS JO, Systems Architect for The New York Times Digital, has architected and/or led development at DeutscheBank, Citibank, Standard & Poor, and ADP. If file system FS2 is mounted on the dir1 directory, when a user changes into that directory and lists the contents, the files from FS2 are shown. Each cylinder group has an inode table, whose relative location is given in the superblock. Backup copies of the superblock can be found in each of the cylinder groups. STEVENS, Software Architect for Hartford Financial Services, is a columnist for Developer.com. Fragments and

blocks that are used to store superblocks, inode tables, group descriptors, and the cylinder group summary area are considered allocated even though they are not allocated to a file. If you are looking for a specific file, you might need to reference several file systems before you find the file because different directories could have been on different
volumes. This can store 16,384 bytes of the file, and the remaining 3,616 bytes are stored in fragments 610 to 613. The $FILE_NAME attribute contains the file name, size, and temporal information. In this case, the compression unit is not compressed, and a run is made for the original data. This process eventually brings us to where the previous
directory entry would have pointed us. Linux always allocates directory entries on block boundaries, so a more general search would examine the first bytes for any file name, not only '.'. You also can use the -l flag to list the allocation status instead of outputting the actual contents. 0000064: 4175 6720 3132 2c20 3230 3034 0000 0000 Aug 12,
2004.... Here is an example: # jls f linux-ext3 ext3-6.dd JBlk Descriptrion 0: Superblock (seq: 0) 1: Unallocated Descriptor Block (seq: 41012) 2: Unallocated FS Block 98313 3: Unallocated FS Block 1376258 [REMOVED] If we are interested in file system block 98,313, we can view the contents of journal block 2 using jcat. Typically, a run contains
the starting cluster location and the size, but a sparse run contains only the size and not a starting location. Unused entries are skipped over by increasing the pointer of the previous entry. It has an inode value of 16,549. The first five clusters are non-zero, the next three clusters contain zeros, and the last four clusters are non-zero. To locate the
group descriptor for a specific group, we need to use the offset location given in the superblock and add it to the group's base address. This scenario is shown in Figure 11.20(C). Like everything in NTFS, the MFT is a file. 11 $Extend A directory that contains files for optional extensions. This node is smaller, so we go to the right-hand child and
compare its value, which is 6. An IA32 (i.e., x86/i386) BSD system will have one DOS partition and one or more BSD partitions inside of it. This section describes where UFS stores file and directory content and how to analyze it. Executable files sometimes contain the names of the files that they open, but our search is unsuccessful. In this case, only a
$INDEX_ROOT is allocated, and it contains three index entry data structures and the empty entry at the end of the list. (If we were using a UFS1 file system, we would also need to calculate the base address for the group.) We read the inode table from block 65,608 and process entry 37, which is the relative location of inode 16,549. The inode for a
file is allocated in the same cylinder group as the parent directory, if space exists. NTFS Documentation. An encrypted file or directory has a special flag set in the $STANDARD_INFORMATION attribute, and each attribute that is encrypted will have a special flag set in its attribute header. The number below each block is the "address" of the entry,
which is there to make it easier to describe each scenario. Boston: Houghton Mifflin, 2000. Here is the output of running fls on a directory in our file system image. For example, fls is a file name category tool (the f) that lists (the ls), and the istat tool is in the metadata category (the i) that displays statistics (the stat). Byte 4 shows the namespace as 1,
which means it is a user attribute. Each index record is given an address, starting with 0. Now let's make this complex by looking at how values are added and deleted. For example, a hat can be stored in a short-round box, and a poster can be stored in a long-round box. # fls -f linux-ext3 ext3-8.dd 32577 r/r 32578: only_live_twice.mp3 r/r 32582:
goldfinger.mp3 r/r 32580: lic_to_kill.mp3 r/r 32581: diamonds_forever.mp3 This might look like an innocent directory of James Bond MP3 files, but notice that the inode numbers and what we know about inode and directory entry allocation. "Windows NT 4.0 and Windows 2000 OEM Support Tools." February 2, 2000. NTFS uses the file reference
address to refer to MFT entries because the sequence number makes it easier to determine when the file system is in a corrupt state. The time values are set to the current time and the link value is set to 1. "The FAT File System." September 20, 2002. The third column is the inode address. A leaf node is one that has no links from it. The OS does not
initialize a block of inode entries in the inode table until it is needed. "Getting Physical with the Digital Investigation Process." International Journal of Digital Evidence, Fall 2003. We can see this relationship in Figure 14.12 where the directory entry in group 113 points to the inode in block group 2. Therefore, there will not be any data from a deleted
file in the slack space of an allocated fragment. Remember that TSK uses sector addresses instead of cluster addresses. This value can be configured in the file system and is set in the superblock. To locate the content, the addresses given in the 12 direct block pointers are used. After a hard link is created, you will not be able to tell if it is the original
name or a link. The attribute contains a list of data decryption fields (DDF) and data recovery fields (DRF). The tree uses directory entries so it can be processed as a normal directory. Using the number of inodes per group, we determine that inode 2 is in cylinder group 0. The fourth unit is sparse, and the corresponding data are all zeros. One is a
size value, and the other two are block addresses where the attributes are stored. Non-resident attributes are stored in cluster runs, which are consecutive clusters, and the run is documented using the starting cluster address and run length. "Detailed Explanation of FAT Boot Sector." Microsoft Knowledge Base Article Q140418, December 6, 2003e.
The time values are stored in UTC, so your analysis tool will need to convert that to the time zone where the computer was actually located so that the displayed time values are correct. 224 $EA Used for backward compatibility with OS/2 applications (HPFS). Many of the fields are non-essential, and in this section I will focus on the essential data and
the non-essential data that could contain evidence. It can be seen in Figure 11.20(B). Report DMCA In A Practical Guide to Enterprise Architecture, six leading experts present indispensable technical, process, and business insight into every aspect of enterprise architecture. Note that the results are more difficult to interpret when there were file
creations in between the file deletions. 0000016: 7777 7777 2e64 6967 6974 616c 2d65 7669 wwww.digital-evi 0000032: 6465 6e63 652e 6f72 6700 0000 0000 0000 dence.org....... We see that there are seven bytes of padding in the content, the name length is six bytes, and that the name is "source." The name ends in byte 12, so the next 8-byte
boundary is byte 16. Further, Solaris uses a new cylinder group after allocating the first 12 blocks of the file, which prevents carving tools from working. An example of the block and fragment relationship is shown in Figure 16.6 where block 64, which contains fragments 64 to 71, is fully allocated to a file, and fragments 74 and 75 in block 72 are
allocated to a file. Volume System Tools The contents of a disk are organized into volumes, and TSK includes one tool that will list the partition layout of a volume. The file size is used to determine how many fragments are being used. This can make carving more difficult. If we return to our previous example, this attribute's run shows that VCN
addresses 0 to 4 map to LCN addresses 48 to 52, VCN addresses 5 to 6 map to LCN addresses 80 to 81, and VCN addresses 7 to 10 map to LCN addresses 56 to 59. The uploader already confirmed that they had the permission to publish it. The UFS1 file system has administrative data with a cycle of 3 and the UFS2 file system has the data at a
constant offset for each group. The group descriptor also tells us that the last allocated inode entry was 16,650. The indexing process makes a tree of the strings in an image so that you can more quickly find the occurrences of specific strings. 2 $LogFile Contains the journal that records the metadata transactions. To convert cluster 9 to its sector
address, we need the sector address of cluster 2, which is 1,632: (Cluster 9 Cluster 2) * 2 (Sectors per Cluster) + Sector 1,632 = Sector 1,646 The fsstat tool in TSK dumps the contents of the FAT structures. /system32/ntio412.sys [REMOVED] Wed Aug 11 2004 19:33:27 2048 mac /bootstat.dat 1024 mac /system32/config/default.LOG 1024 mac
/system32/config/software.LOG Wed Aug 11 2004 19:33:28 262144 ma. Figure 11.20. Every file has a $DATA attribute, which contains the file content. A list of directory entries where two unallocated names are in between two allocated names, and we must advance through the unused space to find the unallocated names. The new entries would
have the fields given in Table 14.2. Table 14.2. Directory entry values after creating a new file. The user can assign the name and value pairs. This section describes where UFS stores the data and how we can analyze them. The listing is shown here: # fls -f linux-ext3 ext3-8.dd 1840555 r/r 1840556: log-001.dat r/r 1840560: log-002.dat r/r 1840566:
log-003.dat r/r 1840569: log-004.dat r/r 32579: snifferlog-1.dat r/r 1840579: log-005.dat r/r 1840585: log-006.dat To find the executable that could have created this file, we search for the full path of the file. Cylinder groups are so named because they were aligned on cylinder boundaries. The tools currently require a raw partition image as input, but
version 2.00 will support disk images. This area will be expanded in the 2.00 release. The UFS block and fragment bitmaps are discussed in more detail in Chapter 17. We have found our value with only three comparisons. We also can see that passwords.txt was deleted before sniffer was and that random.dat was deleted after mytools.zip. The
addresses of all allocated data units also will be shown. The last allocated pointer is also updated. Consider the previous analogy that described an MFT entry as a large box that is initially empty. The index nodes can be stored in two types of MFT entry attributes. For example, a directory index entry contains a few header values and a $FILE_NAME
attribute. Figure 11.15. For example, we defined the logical file system address as the address assigned to file system data units and the logical file address as relative to the start of a file. The header is generic and standard to all attributes. The group descriptor turns out to be in block 65,600, and it shows us that the inode bitmap is located 168
bytes into the group descriptor. It has been reported that if the administrator, domain controller, or other account that is configured as the recovery agent is compromised, any file can be decrypted because that account has access to all files [Microsoft 1999]. AMBLER, senior consultant with Ronin International, specializes in O-O analysis/design,
agile modeling, and architectural audits of mission-critical systems. What makes this confusing is that the MFT has an entry for itself. . A $LOGGED_UTILITY_STREAM attribute is created for the file, and it contains the keys needed to decrypt the data. Attributes are similar to smaller boxes inside the larger box where the smaller boxes can be any
shape that most efficiently stores the object. Therefore, we will look for directories in block group 2 that could have been the parent directory for snifferlog-1.dat. JAMES McGOVERN, Enterprise Architect for Hartford Financial Services, is a co-author of the bestselling books Java Web Services Architecture and Agile Enterprise Architecture. The
second run is the same size as a compression unit, and it is sparse, so we know that there are 16 clusters of zeros. This inode is manually processed in Chapter 17 and is included here as a reference about what data exist. A tree is a group of data structures called nodes that are linked together such that there is a head node and it branches out to the
other nodes. In theory, any attribute could be encrypted, but Windows allows only $DATA attributes to be encrypted. Figure 11.1. An MFT entry has a small header, and the rest of it is used to store different types of attributes. Consider an intrusion investigation where there is an Intrusion Detection System (IDS) that contains log entries
corresponding to the attack. If the bit is a 0, the block or fragment is being used. The OS will look for consecutive fragments that have enough space for the data. Let's look at a simple example to examine each of these scenarios. Each cylinder group has 32,776 fragments and 8,256 inodes. This can take a very long time for large file systems. While
this works, most of us do not copy documents this way because we can remember entire words, and it is more efficient to transfer one or more words at a time. There are three situations that can occur with each compression unit: All the clusters contain zeros, in which case a run of sparse data is made for the size of the compression unit and no disk
space is allocated. In some cases, we might want to search the file system for blocks that were previously used by a directory. There will be two entries with the values given in Table 14.1. Table 14.1. Directory entry values for a new directory. The A-time for the root directory is updated. A fragment is a group of consecutive sectors, and a block is a
group of consecutive fragments. The same file types and permissions are used with UFS as we saw with ExtX. An attribute also has an identifier value assigned to it that is unique to that MFT entry. The file size field is 4 bytes and, therefore, the maximum file size is 4GB. As we will see in the later chapters, this method of saving the time is much
different from other file systems, which save the time as the number of seconds since a given time. London: Academic Press, 2004. Page 26 Brenner, Susan, Brian Carrier, and Jef Henninger. If we think that the IDS was compromised, we should acquire it at the disk level so that we can analyze all the data. Page 18 The file system category of data
includes the general layout and size information of the file system. . If a file system is using hash trees, then a compatible feature flag will be set in the superblock. This has historically involved taking the rotational delay of the disk into consideration to find the optimal location for the next block, but now the next consecutive block is used. "Keeping
an Eye on Your NTFS Drives: The Windows 2000 Change Journal Explained." Microsoft Systems Journal, September 1999. The groups include disk tools, volume tools, file system tools, and searching tools. "Encrypting File System in Windows XP and Windows Server 2003." 2002. Page 9 One of the most important concepts in understanding the design
of NTFS is that important data are allocated to files. Their bits are set to 0, and the proper bookkeeping values are updated. The long file name version is discussed in the next section of this chapter. The amount of padding for the name can be calculated using the name length and the amount of padding for the value is given in byte 5. Page 19 The
content category includes the file and directory content data. It also has flags to identify if the value is compressed or encrypted. We can see this in Figure 11.6. Figure 11.6. Example runlist with three runs of allocated clusters. You can think of an MFT entry as a large box that is used to store your possessions. "Inside Win2K NTFS." Part 2, Windows
and .Net Magazine Network, Winter 2000. To determine all the clusters in this file, we will need to refer to the FAT. Further, the bitwise mask and shift values are also given so that a byte address can be converted to its block address and vice versa. Page 24 Bates, Jim. Linux will not allow an entry to cross a block boundary. Solomon, David , and
Mark Russinovich . The first field in each MFT entry is the signature, and a standard entry will have the ASCII string "FILE." If an error is found in the entry, it may have the string "BAAD." There is also a flag field that identifies if the entry is being used and if the entry is for a directory. For example, Acme Software could decide that its OS would be
faster if an area of the file system were reserved for an address book. Bejtlich, Richard. The hash trees in ExtX are similar to the B-trees that were discussed in Chapter 11, "NTFS Concepts," so refer to that section for an overview of how trees are used in a directory. The A-time of the root directory is updated. The first two entries are for '.' and '..'
and the last entry points to the end of the allocated block. Node A contains three values and four children. "Volume Serial Numbers & Format Verification Date/Time." Digital Detective White Paper, October 2003. Bytes 0 to 3 show the record length as 48 bytes (0x30). The MFT entry and sequence number are combined, with the sequence number in
the upper 16-bits, to form a 64-bit file reference address, as is shown in Figure 11.3. Figure 11.3. Example of the MFT entry address and sequence number combining to form a file reference address. The first run is the same size as a compression unit, so we know it is not compressed. There also could be unused space at the end of the block allocated
to a superblock, cylinder group summary area, or group descriptor. The valid range is from 0 to 127, which gives a year range of 1980 to 2107. We will cover each of the file system metadata files in Chapter 12, but they are listed in Table 11.1 as an easy reference. The directory entry is unallocated by adding its record length to the record length field
in the previous directory entry, which belongs to the 12.jpg file. Landis, Hale, "How It Works: DOS Floppy Disk Boot Sector." May 6, 2002. As we will later see, this information also exists in the group descriptor of each group. New blocks are added as needed, and they are wiped before use. Byte 45 shows the tenths of a second for the create time,
which is 163 (0xa3). 8 $BadClus Contains the clusters that have bad sectors. It was simply moved because fff.txt was deleted. Therefore, recovery could try to locate an indirect block pointer and reconstruct the file. For example, security information is stored in an index, as is quota information. Fortunately, there are many tools that will convert these
values for you so that you do not have do always do it by hand.[1] Many hex editors will show the date if you highlight the value and have the correct options set. The mmls was used in Chapters 5, "PC-based Partitions," and 6, "Server-based Partitions," of this book, and it supports DOS (dos), Apple (mac), BSD (bsd), Sun (sun), and GPT (gpt)
partitions. Source of a Moved File While investigating a Linux system that has been compromised, we come across a file called snifferlog-1.dat that contains network packets. The five available fragments are found using the bitmap (or one of the bookkeeping lists), and they are located in 74,242 to 74,246. The space we took had been used by a
deleted file. "Windows 2000 Server Operations Guide (Part 1)." n.d. . If the content is over roughly 700 bytes in size, it becomes non-resident and is saved in external clusters. Therefore, it is unlikely that it was full when the file was created unless there were a lot of files that were deleted since then. Overview ExtX has several methods for assigning
names to a file or directory, and this section examines three of them. Consider Figure 11.13(A), where we see node A on top and it links to nodes B and C. The istat tool in TSK shows the details of a directory entry and its output for this entry is given next. The attribute header flag identifies whether it is compressed, and the flags in the
$STANDARD_INFORMATION and $FILE_NAME attribute also show if the file contains compressed attributes. A technique that is commonly used in the preservation of a system is to make duplicate copies of the hard disks so that they can be brought to a lab for a dead analysis. We read the dir1 contents from block 66,816 and process the contents as
a list of directory entries. It might be possible to infer about the order in which files were deleted by using the pointer values. We can see this process in Figure 11.11. Each directory entry has eight bytes of static fields in addition to the name, so the minimum record length can be determined by adding eight to the name length and rounding up to a
multiple of four. We can get the details about a specific metadata entry by using the istat tool. We know from the allocation strategy of Linux that files are typically allocated an inode in the same block group as the parent directory. Examples of A) a tree with 5 nodes and B) the same tree that is sorted by the node values. Most forensic tools do not
process the contents of a file system journal. Paul Bakker has been working on adding indexed searches to TSK and Autopsy, and that feature will be part of the 2.00 release (. The -r flag will cause the tool to recursively analyze directories, and the -l flag will look up the metadata and list the temporal data along with the file name. The content of this
attribute is a large buffer that contains one or more index records. Table 10.6. Flag values for the directory entry attributes field. 32 $ATTRIBUTE_LIST List where other attributes for file can be found. "Windows XP Professional Resource Kit Documentation." Chapter 13File Systems, 2004. Page 8 We saw in Chapter 1, "Digital Investigation
Foundations," that the first phase of a digital investigation is the preservation of the digital crime scene. When a file is moved, the new file has an updated C-time, but the M- and A-times remain the same. Newer Sparc Solaris systems might use an EFI partition table instead of a VTOC. There is no easy method for determining if an inode has been
reallocated since the file name was deleted. Instead, a special run is created for the zero clusters. Disk Tools There is only one disk tool in TSK, which is the diskstat tool. Incident Response and Computer Forensics. This output merges the data from the superblock, the cylinder group summary area, and the group descriptors. Scalability is provided by
the use of generic data structures that wrap around data structures with specific content. Figure 16.4. An example layout of a UFS1 cylinder group. In fact, the files that contain the administrative data can be located anywhere in the volume, like a normal file can. We can see this in Figure 14.6(A), where we have three files in a directory. One example
of a generic wrapper is that every byte of data in an NTFS file system is allocated to a file. A directory listing would skip over it, but the data still exists. A more detailed description can be found in "The Sleuth Kit Informer, Issue 16" [Bakker 2004]. When a file is copied, the A-time on the source file is updated, and the destination file has all new M-, A-
, and C-times. We are curious about how the files were deleted, though, and we parse the directory entries to determine the values given in Table 14.3. The names are listed in the order in which they exist in the directory. This would allow us to recover deleted files in each partition, but we would not be able to analyze the sectors that are not
allocated to partitions. The $STANDARD_INFORMATION attribute contains temporal, ownership, and security information. The name of each file system metadata file begins with a "$," and the first letter is capitalized. The superblock references an area of the file system called the cylinder group summary area, which gives a summary of the cylinder
group usage. We will examine the $ATTRIBUTE_LIST attribute in the "Metadata Category" section of Chapter 12. Figure 10.3. Breakdown of the time value and the conversion of 10:31:44 a.m. to its FAT time format. With UFS1, all inodes are initialized when the file system is created. If there are children nodes, their index record addresses are given
in the index entry. The third unit compresses to 10 clusters, so the compressed data is written to disk in a run of 10 clusters, and a sparse run of six clusters is added to account for the compressed data. Blocks are allocated for each of these chunks using a first-available strategy. Many file systems now support journaling because it saves time when
booting large systems. Page 17 Carrier, Brian. . The allocation status of an MFT entry also can be determined from the $BITMAP attribute in the $MFT file, which is shown in Chapter 13. The index entries are organized into nodes of the tree and stored in a list. This directory has three index entries in its resident $INDEX_ROOT attribute and three
index records in its non-resident $INDEX_ALLOCATION attribute. Page 20 The metadata category of data contains the descriptive data for a file or directory. Microsoft calls these files metadata files, but this may cause confusion because we also refer to file metadata. Encryption is a process that uses a cryptographic algorithm and a key to transform
plaintext data to ciphertext data. The sequence number can be useful during an investigation, but in this chapter I will primarily refer to the file number, or MFT entry address, for simplicity. The data structures of both types of directory entries are given in the next chapter. Therefore, the data structure has a field that identifies how long the name is
and where the next directory entry can be found. Extended attributes are a name value pair, and there are currently two "types": user name space and system name space. Cylinder Group Descriptor As previously stated, the file system is organized into cylinder groups. When we locate where jjj.txt should fit, we identify that it should be at the end of
node C, following the iii.txt name. We will later see how this can be used to detect when an inode has been reallocated since a file name was deleted. The first group starts at the beginning of the file system, and the number of fragments in each group is given in the superblock. Analyzing the executable may shed light on those answers. If the file
system is UFS2, the boot code might also occupy additional sectors. Figure 11.14 shows a B-tree with names as values instead of numbers. The M-time and C-time are updated to reflect the inode changes. "Inside Encrypting File System." Part 2, Windows and .Net Magazine Network, July 1999. Typically, the content category includes equal-sized data
units that are allocated for files and directories. File Allocation Example The high-level process for creating the /dir1/file1.dat file is to locate the dir1 directory, create a directory entry, allocate an inode, and then allocate blocks for the file content. This section gives an overview of how it is done and what exists on disk. A non-resident attribute stores
its content in an external cluster in the file system. The length of the entry is rounded up to a multiple of four. Decryption is a process that uses a crypgraphic algorithm and a key to transform ciphertext data to plaintext data. 5 . The bottom of the istat output lists each fragment that the file has allocated. An index record has a static size, typically
4,096 bytes, and it contains a list of index entries. Each cylinder group has a group descriptor data structure that provides details the group. After merging the fragmented runs, we see that there is one run of content, one sparse run, more content, and another sparse run. Notice that there are two superblocks in the first group. Note that the volume
label in the boot sector was set to "NO NAME." The third and fourth lines are for a second directory entry, and we see that the name of this file is "RESUME-1.RTF." The attribute value at byte 43 is 0000 0010 (0x20), which means that only the archive attribute bit is set. Technically, any file that an OS or an application creates could be designed as a
feature in a file system. We do not know the order in which allfiles.tar, config.dat, and delete-files.sh were deleted, but we do know they were before the other files and they are the first three files when the names are sorted. Next, we need to locate the inode table by using the starting fragment of the group and the inode table offset, which is given in
the superblock. The combination of no official specification and one dominant application that creates the file system makes it difficult to differentiate between the application-specific properties and the general properties of the file system. The last 5 bits are for the hour and have a valid range of 0 to 23. Lastly, we locate the specific inode in the table
by subtracting the inode address from the address of the first inode in the group. The six combinations of the relative order that three consecutive directory entries could be unallocated. If you are author/publisher or own the copyright of this documents, please report to us by using this DMCA report form. # hfind NSRLFile.txt
FBF4C1B7ECC0DB33515B00DB987C0474EC3F4B62 FBF4C1B7ECC0DB33515B00DB987C0474EC3F4B62 MOVELIT.GIF Searching Tools The last major category of tools in TSK is searching tools. When stored as a normal attribute, one run of length 12 may be created for the file, as shown in Figure 11.8(A). For example, if we have an unallocated
data structure with a file reference number in it, we can determine if the MFT entry has been reallocated since this data structure used it. If we spend some more time determining the relative order of these deletions, we conclude that the files may have been deleted in alphabetical order. This was discussed in Chapters 8 and 14. . Following them are
entries for every file and subdirectory in the directory. A sparse run follows the compressed run to make the total run length equal to the number of clusters in a compression unit. . An example of converting the time 10:31:44 a.m. to the FAT format can be found in Figure 10.3(B). Here you will find the data unit addresses that a file has allocated, the
size of the file, and temporal information. If someone is shown the ciphertext data, they should not be able to determine the plaintext data without knowing the key. The allocation status of a directory entry structure is determined based on whether it is pointed to by an allocated entry. Now consider an index with 15 entries, which do not fit into an
$INDEX_ROOT but fit into an $INDEX_ALLOCATION attribute with one index record. Also the network packets in only_live_twice.mp3 have timestamps before the packets in snifferlog-1.dat, and the timestamps in lic_to_kill.mp3 are after the snifferlog-1.dat times. Inode 1 used to be used for bad blocks. We find space in between two allocated names,
and the M- and C-times of the directory are updated. Inodes are allocated on a first-available basis in the block group. The chunks of data that are transferred each time are typically a multiple of 512 bytes, because that is the size of most disk sectors. To find the ending location of the value, we subtract the starting byte from the record length and the
padding length (48167 = 25). This output is normally processed using the mactime tool to make timelines, so it is not user friendly. With these values, we can determine the layout of the file system. We do this by reading the blocks that are listed in the inode and processing each entry. "Analysis of Reported Vulnerability in the Windows 2000
Encrypting File System (EFS)." 1999. Encrypted Attributes NTFS provides the capability for attribute contents to be encrypted. If it is, it is processed, and if not, we advance four more bytes and test that location. Fortunately, only the minimum number of fragments are allocated and wiped. Many of these will be discussed in detail in Chapter 12. The
results and specifications can be found on their Web site (. New York: Wiley Publishing, 1995. To make things confusing, it is possible for a directory to have a $DATA attribute in addition to the $INDEX_ROOT attribute. We will do this with the ils tool in TSK. Every file and directory has at least one entry in the table, and the entries by themselves are
very simple. If the file is larger, the indirect blocks must be read and their contents processed as a list of block addresses. Base MFT Entries A file can have up to 65,536 attributes (because of the 16-bit identifier), so it may need more than one MFT entry to store all the attribute headers (even non-resident attributes need their header to be in the
MFT entry). "Description of NTFS Date and Time Stamps for Files and Folders." Microsoft Knowledge Base Article 299648, 2003. It could be possible for data to be hidden in directory entry structures. VIKAS SHARAN is managing partner of Lozoic and architecture team member at Baypackets. Note that this is different from ExtX, which started with
inode 1. . We will see an example of this analysis procedure in the following scenario section. In Microsoft's implementation of NTFS, the MFT starts as small as possible and expands when more entries are needed. . Figure 16.2 shows how a group descriptor could be laid out. We look at the fsstat output and determine that the log directory is in block
group 113, which has 99% of its inodes and 48% of its blocks free. The first entry in the table is named $MFT, and it describes the on-disk location of the MFT. "FILETIME." 2004. The rm tool deletes files, and my tests show that the command deletes files in alphabetical order. The fsck tool in Linux can repackage directories to make them smaller and
remove the unused space. Now we are going to look at the basics of some of the standard attributes. This is why the fls output in TSK gives both the directory entry and inode types. The inode address of our sniffer log was 32,579, which could have been allocated in between only_live_twice.mp3 and lic_to_kill.mp3. Extended Attributes Extended
attributes were added to UFS2, and they provide an additional location to store descriptive data about a file or directory. The first block of the directory is the root node, and it contains the '.' and '..' directory entries. This is analogous to copying a document by hand and reading a letter, punctuation mark, or space from the original and writing it to
the duplicate. Analysis Techniques Analysis of the file name category of data involves listing the names in a directory so that we can find a specific files or files that have a given pattern. If the file were moved within the same block group or to a different file system, then this technique would not work. If the size of the file is unknown or if it slowly
increases, blocks will be added one by one using a next-available strategy. To locate the inode bitmap, we must first locate the group descriptor, which is 48 fragments from the start of the group. An example from our UFS1 image is shown here: # dstat -f openbsd openbsd.dd 288 Fragment: 288 Allocated Group: 0 Allocation Algorithms The designers
of UFS have spent a lot of time researching allocation strategies and have focused on making block allocation efficient. This process repeats until the end of the block. For the final data in the file, fragments will be allocated instead of a full block. An OS that did not know about hash trees would process the entries in all the blocks without knowing
that they were in a sorted order. Typically, the number of children that a node has is based on how many values each node can store. The latter part of the group descriptor contains bitmaps for the inodes, blocks, and fragments in the group. The dynamic UFS2 inode tables also allow another area for data hiding. The allocation status of a block or
fragment is determined using one of two bitmaps. Similarly, if we want to find the metadata entry that a specific file name points to, we can use ifind with the -n flag. The value also is padded so that the next entry starts on an 8-byte boundary. We can see that consecutive clusters were allocated to this file until we get to entry 17 at bytes 68 to 71,
which has an end-of-file marker (0x0fff ffff). We also can use it when recovering deleted content. If we encrypt it for every user, we waste a lot of space. "How to Cause ScanDisk for Windows to Retest Bad Clusters." Microsoft Knowledge Base Article127055, December 16, 2004b. Figure 14.11 shows the six possible combinations of how three
consecutive files could be deleted. The major difference between hash and B-trees is that hash trees sort the files based on a hash of the file name and not based on the name itself. As we saw with ExtX, the touch command can be used to modify the M- and A-times of any file. We check the status of inode 16,651 and find it to be unallocated. This entry
has three attributes. When compressed, the resulting data needs the same number of clusters for storage (i.e., the data did not compress much). 3 $Volume Contains the volume information such as the label, identifier, and version. A special tool would likely be needed to find a block full of 32-bit or 64-bit addresses. We can verify that we have the
correct number of clusters by comparing the file size with the allocated space. The file system type must be specified with the -f flag and one of the types given previously in parentheses. The OS hides a directory entry by increasing the record length of the previous directory entry so that it points to the entry after the one being hidden. Node B links
to node D and E. Wilson, Craig . Figure 16.5(A) shows an example IA32 FreeBSD disk and Figure 16.5(B) shows an example Sparc Solaris disk. In Unix, directories can be used for both storing files and volume mount points, as we discussed in Chapter 4. For example, in the (A) scenario entry 1 was deleted first, and the length of entry 0 was increased
to point entry 2. Therefore, the inode table for inode 2 will start in block 56. When a deleted file name is found, care must be taken when analyzing it. Figure 11.4. Our example MFT entry with the header and content locations specified. The types of data in this category vary depending on the file system type. . The most common use of asymmetric
encryption is where one of the keys is made public, such as "spot," and the other is kept private, such as "felix." Anyone can encrypt data with the public key, but it can be decrypted with only the private key. The base address for a UFS2 group is the start of the group, but it staggers for a UFS1 cylinder group. In the following example, we find that
NTFS cluster 3,456 has been allocated by the $DATA attribute of MFT entry 18,080. "File Deletion in MS FAT Systems." September 23, 2002. Type Identifier Name Description 16 $STANDARD_INFORMATION General information, such as flags; the last accessed, written, and created times; and the owner and security ID. This is done by dividing the
block address by the number of fragments in a group. Cybercrime: The Investigation, Prosecution, and Defense of a Computer-Related Crime. In fact, there is a small flaw in the NTFS design because it creates a temporary file named EFS0.TMP and it contains the plaintext version of the file being encrypted. We are looking the entry for the file1.dat
file. An index in NTFS is a collection of attributes that is stored in a sorted order. This has one value in the root node and two children nodes. Bits 5 to 8 are for the month, and the valid values are 1 to 12. Inside the group descriptor are bitmaps and the standard data structure fields. The group descriptor is allocated a full block and has a combination
of standard fields and a wide-open area that can be used for various tables. Recall that only the block address is given in the inode, but istat will give the addresses of all fragments in the block. We are about to move on to another part of the system when we notice the odd sequence of inode addresses in the listing. We next look at hard and soft links
and then hash trees. 208 $EA_INFORMATION Used for backward compatibility with OS/2 applications (HPFS). The file that allocated entry 313 is deleted, and the entry is reallocated to a new file. Any remnant data in nodes A and C from previously deleted files may now be gone. Page 25 TSK contains over 20 command line tools, which are organized
into groups. Figure 14.8. Example where a directory in FS1 contains three files, but when FS2 is mounted on the directory, they are not seen. The inode contents show us that the contents of dir1 are located in block 66,816. I will refer to these special files as file system metadata files. Unencrypted copies of file content might also exist in unallocated
space if only some directories and files were encrypted. If we want to view unallocated file names as well, we ignore the reported size of the entry and calculate how long the entry should be and advance to that point. We move eee.txt from node I to node F and move bbb.txt from node B to node I. Each cluster was 2 sectors in size, so we can see in the
parentheses that there are 18 sectors in the cluster chain. So we divide it in half and move ggg.txt to the top-most node. See the "File System Category" section in Chapter 12. . A symmetric algorithm uses the same key to encrypt and decrypt data. The sizes of a block and fragment are given in the superblock, and we simply advance from the start of
the file system until we reach the specific address. Redmond: Microsoft Press, 2000. We determine where inode 16,549 is located by dividing it by the number of inodes per group and determine that it is in cylinder group 2. A block is addressed using the address of the first fragment in the block. That sounds easy, but as we will see it results in two
nodes being deleted and five new nodes being created. If the acquisition tool encounters an error while reading data from the suspect drive, many of the tools will write zeros to the destination. The tool displays the total number of sectors and the user-accessible sectors, which show if an HPA exists. . Figure 14.12. Page 6 Brenner, Susan, Brian
Carrier, and Jef Henninger. It is common for the final data in a file to exist in some other part of the cylinder group because it needed only one fragment. For example, consider a file that should occupy 12 clusters. Clifford, Ralph, ed. Figure 11.9. An attribute with two compression units that do not compress, one unit that is sparse, and one unit that
compresses to 10 clusters. A hard link is an additional name for a file or directory in the same file system. For example, there are other methods that could be used to initialize a file system that Microsoft does not use, and it is not clear if they should be considered "valid NTFS" file systems. Having multiple copies of this data also allows them to be
compared if tampering is suspected. All other fragments are for file and directory content. The UFS group descriptor is much larger than its ExtX counter part, although much of it is non-essential data. Figure 14.10 shows an example where we have two unallocated directory entries in between two allocated entries. We read the root directory
contents from block 1,096 and process the contents as a list of directory entries. This area contains a table with an entry for each cylinder group and the entry documents the number of free blocks, fragments, and inodes. The attribute header data structure is given in the "Attribute Header" section of Chapter 13. Figure 16.4 shows an example layout
of a group. The top tree shows 'jjj.txt' added to node C, and the bottom tree is the result of removing node C because each node can have only three names. High-level descriptions of the file system components have been published, but low-level details are sparse. Contains the root directory of the file system. If an entry has more than one attribute of
the same type, this identifier can be used to differentiate between them. If a new node needs to be allocated for the tree, $BITMAP is used to find an available index record; otherwise, more space is added. When compressed, the resulting data uses fewer clusters. We also need to deallocate the one block and five fragments that the file has allocated.
Here is an example output (which has been reduced so that it will fit the width of the book): Wed Aug 11 2004 19:31:58 34528 .a. We also can list the details of several metadata structures by using the ils tool. The non-base entries will have the base entry's address in one of their MFT entry fields. The exact ordering of allocating the different data
structures might vary by OS, and the order presented here might not reflect an actual system. UFS2 has dynamic inodes where they are initialized when they are needed, and the space in the inode table can be used for file content if all other blocks in the file system have been used. . Example Image To close this section, I will run the fsstat tool from
The Sleuth Kit (TSK) on the UFS1 image. If the system were a Solaris system, the inode pointer would also be cleared. Note that the attribute name is different from the type name. With a UFS file system, the first step is to locate the superblock at one of the standard locations. Figure 16.6. Two UFS blocks with 8 fragments in each. If we were looking
for the file ggg.txt, we would look at the values in the root node and determine that the name alphabetically falls in between eee.txt and lll.txt. Every MFT entry also has a 16-bit sequence number that is incremented when the entry is allocated. UFS2 also has extended attributes to store additional descriptive data about a file. 192 $SYMBOLIC_LINK
Soft link information. The OS creates a soft link using a symbolic link, which is a special type of file. When the inode is allocated, its contents are cleared and the M-time, A-time, C-time, and Create-time (for UFS2) are set to the current time. Most file systems exist to read and write file content, but NTFS exists to read and write attributes, one of
which happens to contain file content. An older version has only the name, inode address, and length values. On the other hand, the fragments might allow more data to be recovered than with ExtX, which wiped all bytes in a block. Byte Range Description Essential 03 Record length Yes 44 Namespace (see Table 17.10) No 55 Content padding Yes 66
Name length Yes 7(7 + name length) Name Yes (After name and padded to 8-byte boundary) Value Yes The name is padded so that the value starts on an 8-byte boundary. Microsoft has made changes to the file system with each new release of Windows, and I have noted the differences here. 1,012 12 24 config.dat 20 20 44 readme.txt 104 20 64
allfiles.tar 20 20 84 random.dat 64 20 104 mytools.zip 44 20 124 delete-files.sh 24 24 148 sniffer 876 16 164 passwords.txt 860 860 We can make some general observations from this output. The free block and fragment counts are updated in the group descriptor and cylinder group summary area. The one exception is non-base MFT entries, which
are discussed next. "INFO: Understanding Encrypted Directories." Knowledge Base Article 248723, 2003. 3rd ed. When a new entry is created, the OS examines each existing entry and compares its record length with the name length. The size of the compression unit is given in the attribute header. Page 13 NTFS uses index data structures in many
situations, and this section describes them. This data can be displayed by using the fsstat tool, which will read the boot sector or superblock and other data structures that are specific to the different types of file systems. When looking for a free inode in a cylinder group, a first-available strategy is used. The hash trees still use directory entry data
structures, but they are in a sorted order. The $DATA attribute can store any content that an application or user wants to store there. The first identifies the original location of a file, and the second identifies a possible order in which files were deleted. Also notice that goldfinger.mp3 has a larger inode than the other files, and it is in the middle of the
directory. Directories allocate blocks that will contain a list of directory entry data structures. Microsoft does not typically place the files in this directory into the reserved MFT entries. From processing this, we learn that the block size is 16 KB and fragment size is 2 KB. A B-tree with file names as values. The base increases by d for every group and
returns to the beginning after c groups. The created day is in bytes 48 to 49 and has a value of 0x304a, which is February 10, 2004. The final state from adding the 'jjj.txt' file. . See the "Application Category" section in Chapter 12. Note that when we save only the contents of each sector, we lose data that data recovery specialists may need.
Therefore, the rule of thumb is to acquire data at the lowest layer that we think there will be evidence. We can see this in Figure 11.2 where the boot sector is used to find the first MFT entry, which shows that the MFT is fragmented and goes from clusters 32 to 34 and 56 to 58. The allocation status of any inode is determined using the inode bitmap,
which is located in the group descriptor. diskstat currently runs only on Linux, and it gives the statistics about an ATA hard disk. After the cylinder group has been selected, the OS will allocate blocks based on how many data are being written. It is sometimes useful to examine all the unallocated inode entries because they might contain temporal or
other information from deleted files. It takes the metadata address of the directory as an argument and will list both allocated and unallocated names. Figure 11.13(B) shows the same tree as we saw on the left side, but now with values that are assigned to each node. If the name does not have 8 characters in its name, unused bytes are typically filled
in with the ASCII value for a space, which is 0x20. We will use the group descriptor when determining the allocation status of the group's resources. Both of these attributes are always resident. If no unused areas exist that are large enough for the name, then the name is appended to the list. Therefore, data could be placed in the unused areas
without being noticed until the OS reclaims the area and erases the data. "The Trojan Defense in Cybercrime Cases." Santa Clara Computer and High Technology Law Journal, 21(1), 2004. The top of Figure 11.15 shows this situation, but unfortunately, there are now four names in this node, and it can fit only three. Schneier, Bruce. Therefore, when
the bit is 1, the block or fragment is available. We process the contents of inode 16,651 from the group 2 inode table and decrement the link count by 1 to account for the name being deleted. When additional MFT entries are allocated to a file, the original MFT entry becomes the base MFT entry. A conversion of the date April 1, 2005 to its
hexadecimal format can be found in Figure 10.2(B). File Deletion Example We will now delete the /dir1/file1.dat file using BSD-type methods. Group: 113: Inode Range: 1840545 - 1856832 Block Range: 3702784 - 3735551 Free Inodes: 16271 (99%) Free Blocks: 15728 (48%) We now investigate the sniffer log inode, which is 32,579, in more detail and
determine that it belongs to block group 2. "Tool ReviewWinHex." Journal of Digital Investigation, 1(2), 2004. UFS inodes have 12 direct block pointers, one indirect block pointer, one double indirect block pointer, and one triple indirect block pointer. Figure 16.2. Layout of group descriptor with the standard fields and then bitmaps and other tables
in a wide-open area. As mentioned in the previous section, the order of unallocating the data structures might vary by OS and the order presented here might not be the same as an actual OS. Therefore, a plaintext version of the file exists, and recovery tools might be able to recover the file if its MFT entry has not been reallocated. We can think of
this phase the way we would think of the process of making an exact replica of a building where a physical crime occurred so that investigators can search it for evidence in a lab. . If a directory is using a hash tree, it will have multiple blocks and each will be a node in the tree. The lower 5 bits are for the day of the month, and the valid values are 1 to
31. But it does illustrate how knowing the allocation algorithms can help when these minor details are useful. Digital Evidence and Computer Crime. Therefore, it is best to have a second source of temporal data to help determine the reliability of the time values. The boot code is located in the first sector of the file system and can continue until the
superblock or disk label data structure. One random key is generated for each MFT entry with encrypted data, and it is called the file encryption key (FEK). If the directory entry type is for a directory and the inode has a regular file, it is likely that the inode has been reallocated. Therefore, the full size, content location, and type of file will not be
known. # ifind -f ntfs -d 3456 ntfs10.dd 18080-128-3 Lastly, we can view the contents of any file based on its metadata address instead of its file name using the icat tool. Inode 2 is reserved for the root directory. A third run could start at cluster 56 and have a length of 4. The final result can be found in Figure 11.16. # ils f ntfs e ntfs10.dd
0|a|0|0|1089795287|1089795287|1089795287|100555|1|24755200|0|0 1|a|0|0|1089795287|1089795287|1089795287|100555|1|4096|0|0 [REMOVED] 255|a|256|0|998568000|1100132856|1089795731|100777|1|15360|0|0 256|f|256|0|1100132871|1100132871|1100132871|100777|1|256|0|0 The output was designed so that it can be processed by
another tool, and it is frequently used with the mactime tool to make timelines of file activity. An inode has a last modified, last accessed, and last changed time values, but it does not have a deleted time value. Year: 2,003 Pages: 336 Language: English Topic: 79 Library: mexmat Issue: 29 Identifier: 0131412752,9780131412750 Org File Size:
1,763,827 Extension: chm Top reviews Most recent Top reviews Page 2 UFS2 files and directories can have extended attributes, which are user or system assigned name and value pairs. Both UFS1 and UFS2 superblocks are given with an example file system image. Any pointer with a value of 0 is a sparse block and corresponds to a block of all
zeros. The group descriptor also contains the time of the last write to the group. To do this, we need to be able to locate the inodes, determine their allocation status, and process the extended attributes that might exist. They also support raw and swap images to view individual pages. Microsoft TechNet. This action removes the name from node E
and does not require any other changes. There are four TSK tools in this category, and the names all start with i. In reality, the symbolic link would store the /file1.txt address in its block pointers because the path is shorter than 60 characters. With file systems, such as FAT, that use a list of names in a directory, it is easy to describe why a deleted
name does or does not exist, but with trees it is very difficult to predict the end result. The specific data structures and layout are given in the next chapter. With UFS, the location of the group descriptor, inode table, and backup superblock is specific to each file system, and the offsets are given in the superblock. The $INDEX_ROOT and
$INDEX_ALLOCATION attributes for a directory typically have the name "$I30." Figure 11.7 shows the example MFT entry that we previously used, and its attributes have been given names and types. A user can set the attributes with the setextattr command, and applications can set them with special system calls. Suppose that we used a backup
utility and copied only allocated files. His best sellers include Agile Modeling and The Elements of Java Style. The program then either completes the changes or rolls them back to the original state. Nearly every allocated MFT entry has a $FILE_NAME and a $STANDARD_INFORMATION type attribute. This is an important concept because it explains
why deleted file names are difficult to find in NTFS. If the inode's bit is set to 1, it is allocated. If the file has 12 full blocks, fragments are generally no longer used. If the file system is UFS1, the superblock will be in sector 16, and if the file system is UFS2, the superblock will be in sector 128. We advance ahead using the reported record length of
each entry (we don't care about deleted files) and find the entry whose name is dir1. The UFS2 inode has three fields for the extended attributes. If sniffer had been deleted before mytools.zip, the length of mytools.zip would have been increased to account for it. For example, the key "spot" could be used to encrypt the plaintext into ciphertext, and
the same key could be used to decrypt the ciphertext into plaintext. The dls tool lists the contents of data units, and by default it outputs the contents of all unallocated data units. The updated version uses one of the bytes in the file name length field and uses it to store the file type, such as file, directory, or character device. The first one is because
there is always a superblock in sector 16 (which is block 8 in this file system). This is a dangerous hiding technique, though, because the OS could overwrite it when a new file name is created. 10 $Upcase Contains the uppercase version of every Unicode character. The $INDEX_ROOT attribute is always resident and can store only one node that
contains a small number of index entries. After the OS finishes encrypting the original file, it deletes the temporary file, but the contents are not wiped. "NTFS Undelete (and leap year) Test #1." Digital Forensic Tool Testing, February 2004. The first two fragments in this block are used by another file. The details of the data structures can be found in
Chapter 17. The value is the string "www.digital-evidence.org." Page 3 The FAT directory entry contains the name and metadata for a file or directory. NTFS Implementation When an NTFS $DATA attribute is encrypted, its contents are encrypted with a symmetric algorithm called DESX. Each node contains the files whose hash value is in a given
range. Therefore, an NTFS file system does not have a specific layout like other file systems do. Consider an index with three entries that fit into $INDEX_ROOT. To make things more difficult, consider if fff.txt was deleted. Figure 11.16. If an attribute is resident, the content will immediately follow the header. Solaris has a similar strategy where it
allocates the first 12 blocks (i.e., the direct blocks in the inode) for a file in one cylinder group and then allocates the remaining blocks in other groups. The UFS1 and UFS2 inode data structures are described with file system images in the next chapter. A file will not be deleted until all its hard links are deleted. Cryptography Basics Before we get
into how cryptography is implemented in NTFS, I will give a brief overview of basic cryptographic concepts. Therefore, we know that config.dat was deleted before readme.txt, allfiles.tar was deleted before random.dat, and delete-files.sh was deleted before sniffer. With removable media, it could start in the first sector. This process is shown in Figure
11.12. Figure 10.2. Breakdown of the date value and the conversion of April 1, 2005 to its FAT date format. Every file system metadata file is listed in the root directory, although they are typically hidden from most users. The fls, ils, and icat tools are used to extract the files from the image. Asymmetric encryption uses one key for encryption and a
different key for decryption. The journal records what updates are going to be made to the file system metadata so that a crash can be more quickly recovered from. In fact, it is the only place where the location of the MFT is described; therefore, you need to process it to determine the layout and size of the MFT. The superblock defines a cycle value
c and a delta value d. NTFS is a complex file system and, unfortunately, there is no published specification from Microsoft that describes the on-disk layout. This section covers one of the most common application category features, which is called journaling. The UFS superblock plays the same role as the superblock in ExtX, but the layout and non-
essential data are different. Therefore, we process node C and look at its values. The extended attribute data structure is shown in Chapter 17. Figure 11.18 shows an example of a node in a directory index with four $FILE_NAME index entries. "ScanDisk May Not Fix the Media Descriptor Byte." Knowledge Base Article158869, July 28, 2001. Hash
Trees When the file system is created, the user can choose to use a hash tree to organize the files instead of the unsorted list that I just described. The file has allocated 9 1,024-byte clusters, so there are 9,216 bytes of storage space for the 8,689-byte file. The address in each block pointer is for a full block, except for the final pointer, which could be
for one or more fragments. Page 11 An MFT entry has little internal structure and most if it is used to store attributes, which are data structures that store a specific type of data. To date, all entries have been 1,024 bytes in size, but the exact size is defined in the boot sector. If we are trying to look up a value, we compare it to the root node. NTFS
uses both sparse runs and compressed data to reduce the amount of space needed. 0000032: 5245 5355 4d45 2d31 5254 4620 00a3 347e RESUME-1RTF ..4~ 0000048: 4a30 8830 0000 4a33 7830 0900 f121 0000 .0.0.....0...!.. The jls tool will list the contents of the journal and show which file system blocks are saved in the journal blocks. UFS1 adds
an additional twist because the offsets are relative to a base location that changes for each cylinder group. 0000016: 0000 0000 0000 874d 252b 0000 0000 0000 ....M%+...... Analysis Considerations The file system category of data will not typically include much evidence of an incident. Table 14.3. Values in unallocated directory entry structures for
scenario. The basic information is equivalent to an MFT entry's fixed fields. 144 $INDEX_ROOT Root node of an index tree. The BSD partition locations are defined by the disk label structure, which is located in sector 1 of the DOS partition. We find its entry and observe that it has allocated inode entry 16,651. The M-time and C-time of the inode
32577 directory are the same as those for the goldfinger.mp3 file, which are both after the times for the snifferlog-1.dat file. We view the three allocated directories using fls. A resident attribute stores its content in the MFT entry with the attribute header. 160 $INDEX_ALLOCATION Nodes of an index tree rooted in $INDEX_ROOT attribute. The bit
for this block is set to 0 to show that it is no longer free, and the count of available blocks is decremented. When that value is reached, a new group is selected. A parent node is one that links to other nodes, and a child node is one that is linked to. 2nd ed. If the file name has the value 0xe5 in that byte, 0x05 should be used instead. Its base address is
located several blocks inside the group, and the backup superblock, group descriptor, and inode table all follow in consecutive order. The name is typically in ASCII, but could also use one of the Microsoft code pages if the name uses non-ASCII symbols [Microsoft 2004]. It is also possible for the UFS2 superblock to exist at 256 KB from the start of the
file system, but this is not the default. Showing 1-46 Start your review of A Practical Guide to Enterprise Architecture Mar 20, 2012 Eduardo rated it really liked it Good reference material for those who wish to review or learn the many concepts associated with hardware architectures; although it does not present anything new, it consolidates lots of
information. There are two TSK tools that operate at the file name layer, and their names start with f. See the "Content Category" section in Chapter 12. The first two entries in each directory will always be allocated because they are for the '.' and '..' entries. Analysis Techniques When we analyze the metadata category of data, we are looking for
descriptive data about a file so that we can sort it and obtain more information. Cluster 9 has a 36-byte offset into the FAT32 structure, and we previously calculated that the primary FAT structure starts in sector 38. In this example, all the attributes are resident. We know this because the length of mytools.zip was increased to 44 to account for
delete-files-sh being deleted. Byte Range Description Essential 00 First character of file name in ASCII and allocation status (0xe5 or 0x00 if unallocated) Yes 110 Characters 2 to 11 of file name in ASCII Yes 1111 File Attributes (see Table 10.6) Yes 1212 Reserved No 1313 Created time (tenths of second) No 1415 Created time (hours, minutes,
seconds) No 1617 Created day No 1819 Accessed day No 2021 High 2 bytes of first cluster address (0 for FAT12 and FAT16) Yes 2223 Written time (hours, minutes, seconds) No 2425 Written day No 2627 Low 2 bytes of first cluster address Yes 2831 Size of file (0 for directories) Yes The first byte of the data structure works as the allocation status,
and if it is set to 0xe5 or 0x00, the directory entry is unallocated. Superblock The UFS superblock contains basic information, such as the size of each fragment and the number of fragments in each block. For example, consider a new directory that has only a '.' and '..' entries in a 4,096-byte block. Page 23 To conclude the discussion on UFS, we will
step through an example where we allocate and then delete a /dir1/file1.dat file, which is 25,000 bytes in size, from a UFS2 file system. Before the attribute contents are compressed, the data are broken up into equal sized chunks called compression units. The bitmaps are stored inside of a cylinder group's group descriptor data structure, which we
previously examined. # jcat f linux-ext3 ext3-6.dd 2 Multiple Category There are a few tools that combine the data from the various categories to produce the data sorted in a different order. Each line in the output corresponds to a file being accessed or changed somehow, which we discussed in Chapter 8. Directory entries that have the long file
name attribute set have a different structure because they are storing the long name for a file, and they will be described in the next section. Let's look at the raw contents of two directory entries from the root directory. NTFS uses B-trees, which are similar to the binary tree we just saw, but there can be more than two children per node. Symmetric
encryption is very fast, but it is difficult when sharing the ciphertext data. The OS uses the node descriptors to determine to which block it should jump for a given hash value. 128 $DATA File contents. If a file cannot fit its attributes into one entry, it can use multiple entries. After we fill up the index entries in the index record, we need to add a new
level, and we create a three-node tree. Other OSes could use different allocation methods or even compact the directories to make them smaller. Consider an example file system with 8,192 byte blocks and 1,024 byte fragments. It is here that we will also learn how big each cylinder group is and where the various data structures are inside each
group. Here is that output: # fsstat f fat fat-4.dd [REMOVED] 1642-1645 (4) -> EOF 1646-1663 (18) -> EOF 1664-1681 (18) -> EOF [REMOVED] Here the output shows us the cluster chain for RESUME-1.RTF from sectors 1646 to 1663 and the End of File. The state shown (A) is unique among the other combinations of how those files can be deleted.
The file size is decremented each time a block is deallocated, and it eventually reaches 0. In addition to a number, each attribute type has a name, and it has all capital letters and starts with "$." Some of the default attribute types and their identifiers are given in Table 11.2. Not all these attribute types and identifiers will exist for every file. The last
line of the 'Direct Blocks' section shows that only five of the eight fragments in block 1576 are being used. Exists only in version 1.2 (Windows NT). The base MFT entry will have an $ATTRIBUTE_LIST type attribute that contains a list with each of the file's attributes and the MFT address in which it can be found. 1 $MFTMirr Contains a backup of the
first entries in the MFT. The lower 5 bits are for the second, and it uses two-second intervals. Bookkeeping information exists to make the allocation of files more efficient, such as the location of the last allocated block, fragment, and inode. The standard fields provide bookkeeping information and describe how the latter part of the block is organized.
Figure 14.6. Directory entries contain the file name and inode. To reduce the effect of a hard disk error, the administrative data are staggered so that not every copy of the superblock is on the same platter. Acquisition Tool Testing Acquisition is a crucial part of the investigation process, and the National Institute of Standards and Technology (NIST)
has conducted tests on common acquisition tools. Figure 11.19. The output will show the size and temporal data as well as any permissions fields. # dstat -f linux-ext3 ext3-5.dd 23456 Block: 23456 Not Allocated Group: 2 Lastly, we can view the contents of any data unit using the dcat tool. Instead, it will place a 0 in the block pointer and output a
block of zeros when that part of the file is read. To deallocate the inode, the corresponding bit in the inode bitmap is set to 0, and we update the free inode counts in the group descriptor and cylinder group summary area. Overview Like ExtX, UFS uses inode data structures to store file and directory metadata. # fls -f linux-ext3 a ext3.dd 69457 d/d
69457: . 1996-2004. It exists to make the consistency checking faster. Therefore, we break node C in half, move ggg.txt up a level, and create nodes F and G with the resulting names from node C. Its address is given in the superblock, and it contains a table of usage information for each cylinder group. 6 $Bitmap Contains the allocation status of each
cluster in the file system. If the file's inode has been reallocated, the metadata is no longer relevant to the deleted name. Figure 16.9. Final state after deleting the '/dir1/file1.dat' file. 80 $SECURITY_ DESCRIPTOR The access control and security properties of the file. "FAT Undelete Test #1." Digital Forensic Tool Testing, February 2004a. If the OS is
allocating blocks for a new file and there are available blocks in the cylinder group of the inode, blocks from that cylinder group are used. A user's private key is stored in the Windows registry and encrypted with a symmetric algorithm that uses her login password as the key. If we find a data unit with interesting evidence, we can search all the
metadata entries using the ifind tool with the -d flag. While there is a documented allocation strategy, there is no requirement that an OS must follow it. We read the superblock data structure, which is 2 KB in size and is located 64 KB into the file system. For example, the inode table could start 32 fragments from the start and the group descriptor 16
fragments from the start. This section shows what an index looks like and how it is implemented. The starting location can be calculated using the delta and cycle values in the superblock. "FAT Volume Label Test #1." Digital Forensic Tool Testing, August 2004b. For most cases, an investigator will acquire every sector of a disk, which is what we
cover in this chapter. For each block and fragment, the corresponding bit in the block and fragment bitmaps are set to 1, and the block pointer in the inode is cleared. The group descriptor identifies block 67,896 as the last block allocated. Bits 9 to 15 are for the year, and the value is added to 1980. To determine the bit in the bitmap for a specific
block, we calculate the starting block address of the group and subtract that value from our block address. Analysis Techniques Analysis of the file system category involves locating and processing the general file system information so that additional analysis techniques can be applied. Casey, Eoghan. MICHAEL E. We examine block 67,904 in the
free fragment bitmap and determine that it is not allocated. We will start from the bottom and work our way up. The link count in the inode is incremented by one to account for the new name. Part B shows a soft link to the same file, but this time there is another level of indirection. That is the only action that places the entry in an unallocated status.
4 $AttrDef Contains the attribute information, such as the identifier values, name, and sizes. TSK has tools called jls and jcat that list the contents of some journals. "MS-DOS FORMAT Does Not Preserve Clusters Marked Bad." Knowledge Base Article103548, May 6, 2003c. The allocation strategies can be used when carving data by focusing attention
on only a specific group instead of the entire file system. Each file and directory uses one inode, and it contains the address of the blocks that a file has allocated, the file size, and temporal information. For example, the size of a block is given in both bytes and in fragments. The size of the file is 25,000 bytes, so we need to allocate one block and five
fragments. The content is specific to the type of attribute and can be any size. Analysis Considerations Most of the analysis considerations of ExtX metadata apply to UFS. Using the data in the superblock, we can find the locations of each cylinder group and data block. The contents of a specific journal block can be viewed by using the jcat tool. Slack
space might contain hidden data, though. # istat -f openbsd -z UTC openbsd.dd 3 inode: 3 Allocated Group: 0 uid / gid: 0 / 0 mode: -rw-r--r-- size: 1274880 num of links: 1 Inode Times: Accessed: Tue Aug 3 14:12:56 2004 File Modified: Tue Aug 3 14:13:14 2004 Inode Modified: Tue Aug 3 14:13:14 2004 Direct Blocks: 288 289 290 291 292 293 294 295
296 297 298 299 300 301 302 303 304 305 306 307 308 309 310 311 [REMOVED] 1568 1569 1570 1571 1572 1573 1574 1575 1576 1577 1578 1579 1580 Indirect Blocks: 384 385 386 387 388 389 390 391 We can see that this file is 1,247,880 bytes, so it needs more than 12 of the 8KB blocks. File system journals may turn out to be useful in
investigations, although to date they have not been fully utilized. 192 $REPARSE_POINT Contains data about a reparse point, which is used as a soft link in version 3.0+ (Windows 2000+). 9 $Secure Contains information about the security and access control for the files (Windows 2000 and XP version only). Let us go over some examples. Further, the
name goldfinger.mp3 is 14 characters long, and the name snifferlog-1.dat is 16 characters long, which means that they can use the same sized directory entry. Using the name length, it calculates how long the entry needs to be. The final state of the system can be found in Figure 16.8. Figure 16.8. Final state after adding the 'dir1/file1.dat' file.
Microsoft calls this sequential address the file number. The superblock is located in the beginning of the file system and contains the basic size and configuration information. When a file is deleted, its M- and C-times are updated. On the plus side, because the tools do not show volumes at their mount points, you can see the directory contents of the

mount points. Scenario where the snifferlog-1.dat file was moved from a directory in block group 2. Like FAT, NTFS uses clusters, which are groups of consecutive sectors. Node F becomes empty, and we need to fill it in. We can now view some of this same data with TSK. We can do this by searching the bitmaps for inodes whose bits are 0 and processing inodes that are non-zero. With UFS1 we also will need to calculate what the staggering offset for the group is. SCOTT W. There can be up to three layers of nodes in a hash index tree. You'll find start-to-finish guidance for architecting effective system, software, and service-oriented architectures; using product lines to streamline enterprise software design; leveraging powerful agile modeling techniques; extending the Unified Process to the full software lifecycle; architecting presentation tiers and user experience; and driving the technical direction of the entire enterprise. . The merged data are organized into compression units, and we see that the first two units have no sparse runs and are not compressed. The UFS bitmaps are opposite of normal bitmaps because they are "free" bitmaps. Page 15 TSK contains over 20 command line tools, which are organized into groups. The first two lines show a directory entry with the attribute at byte 11 set to the binary value 0000 1000 (0x08), which is for a volume label. Therefore, snifferlog-1.dat was either originally allocated to a different parent directory and moved to its current location, or it was created in the current directory but the block group was full. Durham: Carolina Academic Press, 2001. If the record length of an entry is larger than it needs to be and the difference is more than the size of the entry being added, the entry is placed in the unused space. This is important for UFS1 where they are located at different offsets for each group. It also shows you which sectors in the volume are not allocated to a partition. There are many types of index entries, but they all have the same standard header fields, which are given in Chapter 13. Cooperstein, Jeffrey , and Richter, Jeffrey . After it was moved, the goldfinger.mp3 file was created, and it overwrote the directory entry and took the next available inode. Most file systems separate the name and metadata, and the name is located inside of the data units allocated to a directory. This is especially true when the hash trees are used because the first block contains a small amount of administrative data, and the rest is unused. Allocation Algorithms The documented allocation method of UFS inodes is the same that we saw for ExtX. Therefore, these files might have been deleted when they were shown in a window sorted by name or by a command, such as rm *. Values for the extended attribute name space field. The other fragments in block 72 can be allocated to other files. When a file is deleted, the record length of the previous entry is incremented so that it points to the entry after the one being deleted. Its contents are shown here: # dcat f fat fat-4.dd 38 | xxd [REMOVED] 0000032: ffff ff0f 0a00 0000 0b00 0000 0c00 0000 ..... Figure 11.10. Here is an Ext3 image with a directory in inode 69457, which contains a deleted file named file two.dat. UFS supports sparse files, which were discussed in Chapter 8, "File System Analysis." If a file has not defined the contents of part of a file or if the block is all zeros, the OS will probably not allocate a block. The space between the last directory entry and the end of the block is unused and could contain data. Trees are useful because they can be used to easily sort and find data. A DDF is created for every user who has access to the file, and it contains the user's Security ID (SID), encryption information, and the FEK encrypted with the user's public key. The inode address is added to the file1.dat directory entry. The file system tools support Ext2/3 (linux-ext2, linux-ext3), FAT (fat, fat12, fat16, fat32), NTFS (ntfs), and UFS1/2 (freebsd, netbsd, openbsd, solaris) file system formats. As we will see in Chapter 5, "PC-based Partitions," a disk that has DOS partitions may not use sectors 1 to 62, and they could contain hidden data. Unused entries are not rearranged to compress the size of a directory. Fortunately, other groups have published what they think the on-disk data structures are [Linux NTFS 2004], and those are included in this book and we use them to dissect a disk by hand. It also contains tables that help find consecutive fragments and blocks of a given size. One hiding technique is to create files in a directory and then mount a volume on the directory so that a casual observer would not notice them. For some systems, our rule of thumb about acquiring at the level where we think there will be evidence means that we need to copy only files. Microsoft calls each entry in the table a file record, but I think calling each entry an MFT entry is simpler and results in fewer terms to remember. To make a hard link, the OS allocates a new directory entry and points it to the original inode. Overview UFS has three types of data structures that store the file system category of data: the superblock, the cylinder group summary, and the group descriptor. In this scenario, we have tied a file back to its original directory using its inode address. With these results, it is likely that the files were not individually deleted one by one and that maybe a script was used or a file manager window. UFS1 and UFS2 use slightly different data structures, but both are over 1 KB in size and contain nearly 100 fields. There are two categories of cryptographic algorithms: symmetric and asymmetric. We also can identify where the inode table and group descriptor for each group are located. To show why we typically acquire at the disk level, we will consider some scenarios. Page 12 The previous section looked at the basic concepts that apply to all NTFS attributes. With UFS, if a file can fill an entire block, it will be allocated a full block. Very large files that need to be recovered, such as a large e-mail file, could be difficult because they might have taken up more than 25% of the cylinder group and the remainder moved to another group. Emeryville: McGraw Hill/Osborne, 2003. We set its bit in the bitmap, update the last allocated inode value in the group descriptor, and decrement the number of free inodes value in the group descriptor and cylinder group summary area. The first unit compresses to 16 clusters, so it is not compressed. Note that this is a file system-level compression and not an external application-level compression that can be achieved by using zip or gzip. It also contains many variations of a single value so that the OS does not have to compute the variation every time. Linux NTFS Project. Each index is given a name and the $INDEX_ROOT, $INDEX_ALLOCATION, and $BITMAP attributes for the index are all assigned the same name in their attribute header. Refer to the "File System Category" sections of Chapters 9, "FAT Concepts and Analysis," 12, "NTFS Analysis," 14, "Ext2 and Ext3 Concepts and Analysis," and 16, "UFS1 and UFS2 Concepts and Analysis," for specific outputs. Notice that the long file name attribute is a bit-wise combination of the first four attributes. The most common usage of an index is in a directory because directories contain $FILE_NAME attributes. Allocation Algorithms When a new file name is created, Linux uses a first-available strategy. After locating the directory content, we process it as a sequence of directory entry data structures. Byte Location Name Record Length Needed Length 12 .. The group descriptor is located using the starting address of the group and offset values from the superblock. You can search for the backup copies using the data structure's signature value. In Windows, a user can choose to encrypt a specific file or a directory. If we use one key for everyone, it is difficult to revoke access from a user without changing the key. On the plus side, blocks should be localized to their inode, which could make file recovery easier, and if deleted data are in a group that has little activity, it could exist for longer than deleted data in other groups. If we can store five values in each node, we can have six children. In this case, we would not be able to recover deleted files, we might not have access to all the temporal data, and we would not be able to find data that has been hidden inside partition or file system data structures. If they are different, it assumes that either it is at the end of a block or that the record length was increased to cover a deleted entry. The reason the data are staggered is to reduce the impact of physical damage to a platter in older disks. Each entry is given an address based on its location in the table, starting with 0. Directories will have a size of 0 and the FAT structure must be used to determine the number of clusters allocated to it. # istat -f fat fat-4.dd 4 Directory Entry: 4 Allocated File Attributes: File, Archive Size: 8689 Name: RESUME-1.RTF Directory Entry Times: Written: Wed Mar 24 06:26:20 2004 Accessed: Thu Apr 8 00:00:00 2004 Created: Tue Feb 10 15:49:40 2004 Sectors: 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 Page 4 Microsoft. The date portion of each timestamp is a 16-bit value that has three parts, which are shown in Figure 10.2(A). Figure 11.2. The relationship between the boot sector and $MFT with respect to determining the layout of the MFT. Adding one file results in removing nodes A and C and adding nodes F, G, H, I, and J. /system32/config/default Another tool that reorders data is the sorter tool, which sorts files based on their content type. If you get to the empty entry at the end of the list, you look at its child. # icat f ntfs ntfs10.dd 18080 File Name Category The file name category of data includes the data that associates a name with a metadata entry. We can use this to get basic information about a group. Page 21 The file name category of data includes the data structures that store the name of each file and directory. General bookkeeping information is also stored in the superblock, such as the total number of free inodes, fragments, and blocks. We previously saw part of the fsstat output when discussing the file system category of data, but the FAT contents were removed. General Acquisition Procedure The general, and intuitive, procedure for acquiring a storage device is to copy one byte from the original storage device (the source) to a destination storage device and repeat the process. The inside of the box is initially empty, but it can be used to store anything as long as it is in a container that is smaller than the box. The final block could be a fragment, and the next fragment might be used by another file. . If hash trees are being used, the file is added to the block that corresponds to the file's hash value. . If so, this block was the first block in a directory. 1 12 .. "UK Computer Misuse ActThe Trojan Virus Defense." Journal of Digital Investigation, 1(2), 2004. This is not typical in a Linux system, and we look at its contents. In this case, the data is compressed and stored in a run on the disk. The mode field and time values are encoded and must be processed before sense can be made of them. This includes the basic file system administrative data that are typically hidden by other file systems. Mandia, Kevin, Chris Prosise, and Matt Pepe. For example, if we are given a block address, we can determine to which group it belongs. We are looking for unused space in the directory. NTFS uses different terms for these addresses. After the changes are made, another entry is made in the journal to show that the changes occurred. Group: 2: Inode Range: 32577 - 48864 Block Range: 65536 - 98303 Free Inodes: 16268 (99%) Free Blocks: 0 (0%) One theory is that it was created in a directory in block group 2 and moved to the directory in block group 113. If the directory is large, $INDEX_ALLOCATION and $BITMAP attributes are also used to store information. Figure 11.4 shows an MFT entry with four header and content pairs. The contents of the indirect block pointers are not cleared, though, so searching for a block of indirect block pointers could help during recovery. If the size is known, the data are broken up into chunks, which are typically the size of 16 blocks. The first step in this process is to locate the root directory, which is easy in ExtX because it is always located in inode 2. When the file grows in size and can fill an entire block, it will be moved to one if needed. The default $DATA attribute that is created when a file is created does not have a name associated with it, but additional $DATA attributes must have one. Larger indexes allocate a non-resident $INDEX_ALLOCATION attribute, which can contain as many nodes as needed. Flag Value (in bits) Description Essential 0000 0001 (0x01) Read only No 0000 0010 (0x02) Hidden file No 0000 0100 (0x04) System file No 0000 1000 (0x08) Volume label Yes 0000 1111 (0x0f) Long file name Yes 0001 0000 (0x10) Directory Yes 0010 0000 (0x20) Archive No The upper two bits of the attribute byte are reserved. Processing an inode is fairly straightforward. aeb/linux/fs/fat/fat.html. Page 22 Some file systems contain data that belongs in the application category. Name Name Length Record Length . The data in inode 2 shows that the directory entry structures for the root directory are located in block 1,096. The maximum MFT address changes as the MFT grows and is determined by dividing the size of $MFT by the size of each entry. At each layer of abstraction, data are lost. This works for only small attributes. We locate inode 139 in the inode table, which is the relative location of inode 16,651, and initialize its values. Newer hard disks do not have the same number of sectors in each cylinder, so this is not an issue, and UFS2 no longer staggers the data. When the entry is reallocated, it has a new sequence number of 2. The time values in UFS1 are 32 bits and 64 bits in UFS2. The $BITMAP attribute is used to manage the allocation status of the index records. The minimum size of a UFS block is 4,096 bytes, and the maximum number of fragments per block is eight. 64 $OBJECT_ID A 16-byte unique identifier for the file or directory. We can see this in Figure 14.8, where part A shows the three pictures in dir1 and part B shows that dir1 now contains the three files in the root directory of volume FS2. "Encodings and Code Pages." Global Development and Computing Portal, 2004a. We still do not know how the file was moved and if it always had that name or if it had an MP3 name. It contains eight deleted file names, and all of the corresponding inodes have been reallocated to new files, so file recovery will not be trivial. The link count is set to 1 for files and 2 for directories to account for the '.' name inside of the directory. One has a bit for each fragment, and the other has a bit for each block, and they should contain nearly the same information, but the block bitmap allows the system to more quickly find consecutive blocks when allocating a large file. Compressed Attributes NTFS allows attributes to be written in a compressed format, although the actual algorithm is not given. Assume that the compression unit size is 16 clusters and we have a $DATA attribute that is 64 clusters in length, as shown in Figure 11.9. We divide the content into four compression units and examine each. For example, config.dat was deleted before readme.txt was deleted because if readme.txt was deleted before config.dat, the length of the config.dat file would have been increased by 20 to cover up readme.txt. The space after the c.txt file is unused. The BSD systems that I tested wrote 0s for the unused bytes in the final fragment of a file, also called slack space. It has the three standard file attributes. To solve this problem, NTFS provides two locations where attribute content can be stored. The table entry for cluster 9 is located in bytes 36 to 39, and we see that the value is 10 (0x0000 000a), which means that cluster 10 is the next cluster in the chain. At the time of this writing, the current version is 1.73, but there are plans for big changes in a 2.00 release. The resulting state is found in Figure 11.17. The block address corresponds to the address of the first fragment in the block. Let us assume that we can fit only three file names per node, and the file jjj.txt is added. To fully understand the layout of UFS, it might be helpful to refer back to Chapter 6, "Server-based Partitions." BSD and Solaris systems have their own partitioning systems. Theoretically, only one of these values needs to exist and the others can be calculated. Chapter 13 shows the data structure for an MFT entry and dissects our example file system image. This might cause a performance improvement, but it is not essential for the file system. The tool runs the file command on each tool and saves the file to a category based on a set of rules. Older hard disks had the same number of sectors per track, which meant that the first sector of every group was on the same platter. r/r 69458: abcdefg.txt r/r * 69459: file two.dat d/d 69460: subdir1 r/r 69461: RSTUVWXY Links and Mount Points ExtX provides both hard and soft links so that users can define multiple names for a file or directory. diskstat was used in Chapter 3, "Hard Disk Acquisition," when we looked for Host Protected Areas (HPA) before acquiring a disk. The superblock contains disk geometry information that was used to most efficiently organize and optimize a file system. The fsstat output is much larger that what is shown here because it details every cylinder group, but the relevant parts are shown. Consider a directory dir1 that is in a file system named FS1. . "Open Source Digital Forensic Tools: The Legal Argument." Fall 2003b. Spafford. Example Image Here is the output from running istat on our example file system image. The FEK is stored in an encrypted state in the $LOGGED_UTILITY_STREAM attribute. Every directory starts off with directory entries for the '.' and '..' directories, which are for the current and parent directory. For example, if an attribute has allocated clusters 48, 49, 50, 51, and 52, it has a run that starts in cluster 48 with a length of 5. Figure 11.17. The next example will extract all unallocated space of the NTFS image: # dls f ntfs ntfs-10.dd > ntfs-10.dls The resulting file will have no structure to it because it simply contains random data units from the file system. This section describes where the data are stored and how to analyze them. Example output of this tool was given in Chapters 9, 12, 14, and 16. File systems that do not contain boot code will not use the sectors before the superblock. The free blocks are reserved so that a system administrator can login and clean up the file system, but an OS might choose to not enforce the limit. One is small and contains only boot code. "Encodings and Code Pages." Global Development and Computing Portal, 2004. If the system crashes, the scanning program reads the journal and locates the entries that were not completed. Table 11.2. List of default MFT entry attribute types. Therefore, the link count for a directory is equal to at least two plus the number of subdirectories it has. One of these entries is allocated for every file and directory, and they are located in the clusters allocated to the file's parent directory. The reason this is more complex is because the layout is not initially organized using compression units. Figure 11.14. The OS starts at the beginning of the directory and examines each directory entry. Nodes C, D, and E are leaves. In TSK, this includes only two tools, which are for the journal in Ext3. The Tao of Network Security Monitoring: Beyond Intrusion Detection. We advance ahead by using the reported record length of each entry (we don't care about deleted files) and find the entry named dir1. When run on an NTFS image, it will show all the file's attributes. MFT Entry Addresses Each MFT entry is sequentially addressed using a 48-bit value, and the first entry has an address of 0. A file that is 20,000 bytes in size will need 2 blocks and 4 fragments. [1] An example is "Decode from Digital Detective" ( ). If the attribute also allocated clusters 80 and 81, it has a second run that starts in cluster 80 with a length of 2. . The attribute type is a numerical identifier based on the data type and we will discuss the default attribute types in the "Standard Attribute Types" section. The American Heritage Dictionary. Analysis Techniques We analyze data in the content category by locating a block or fragment, viewing the contents, and identifying its allocation status. The M- and C-time in an unallocated inode entry might reflect the time that the corresponding file was deleted. Refer to "A Case Study Using dd" in Chapter 3 for a specific example. [2] Microsoft documentation says it reserves only the first 16 entries, but in practice the first entry that is allocated to a user file or directory is entry 24. The third and fifth units have a sparse run and are compressed. When the final data are being written, only enough fragments for the data are allocated. While each type of attribute stores a different type of data, all attributes have two parts: the header and the content. /system32/config/SECURITY 262144 ma. Figure 14.11. Page 14 Cooperstein, Jeffrey , and Richter, Jeffrey . Some of the attributes can be assigned a name and it is stored in UTF-16 Unicode in the attribute header. We find the file name in that node. Page 5 TSK contains over 20 command line tools, which are organized into groups. We also can observe that mytools.zip was deleted after delete-files.sh but before sniffer was deleted. Exists only in versions 3.0+ and after (Windows 2000+). The namespace value can take on one of the values given in Table 17.10. We also learn that the group descriptor is located 40 fragments, and the inode table is located 56 fragments into each cylinder group. Any user that can read the file's contents can read the file's user name space attributes. Many current post-mortem investigation tools do not show volumes at their mount point, and therefore you will need to determine which volume should be there. The type of data in the output of fsstat is different for each file system because different types of data are available. The other files in the directory have names similar to log-001.dat and do not contain network data. "Last Access Date." MSDN Library, February 2003a. We also get a breakdown of each cylinder group with respect to what fragments are in each group and which resources were last allocated. We can see this in Figure 14.9 where we have several files in two leaves. To examine only the allocated names, we process a directory entry structure, jump ahead by its reported size, and process the next entry. The "base" location for each group is calculated using two values from the superblock. Applied Cryptography. If we look at the final states of the six scenarios, only (B) and (D) are the same. If you search the file and find evidence, you can determine from where it originally came by using the dcalc tool. # dcat f fat fat-4.dd 1632 | xxd 0000000: 4641 5420 4449 534b 2020 2008 0000 0000 FAT DISK ..... We can also see that the write time and date are set at bytes 22 to 25 on line 2. There is a lot of bookkeeping information in the superblock, group descriptors, and cylinder group summary area. Of course, the OS that created the file system might not have been following every rule of UFS and did not use these allocation principles. See the "Metadata Category" section in Chapter 12. If the file has a more complex name, there will be a long file name directory entry in addition to a directory entry. File System Metadata Files Because every byte in the volume is allocated to a file, there must exist files that store the file system's administrative data. Analysis Considerations The design and allocation strategies of UFS provide some benefit and challenges for investigators. One block is allocated, and two fragments of the other are allocated to a file. The user's private key is used to decrypt the FEK, and the FEK is used to decrypt the $DATA attribute. The -e flag can be used to output all data units, which is the same as using dd on the image. Bytes 46 to 47 have the created time, 0x7e34, which is 15:49:40. We could have found the value 9 in only two comparisons instead of five if the values were in a list. We need to locate the dir1 directory in the root directory, so we process inode 2. Also on the plus side is that clusters of consecutive blocks are allocated when possible so that fragmentation is reduced, which could help carving tools. When a file has more than one $DATA attribute, the additional attributes are sometimes referred to as alternate data streams (ADS). 256 $LOGGED_UTILITY_STREAM Contains keys and information about encrypted attributes in version 3.0+ (Windows 2000+). We also see that readme.txt was deleted before sniffer because the record length for readme.txt points to sniffer. The node is larger, so we go to the left-hand child and compare its value, which is 5. Refer to Chapter 14 for an example of this. When an attribute is encrypted, only the content is encrypted and the attribute header is not. Note that there is also experimental support for B-Trees in ExtX that are like the NTFS B-Trees, but we do not discuss them in this chapter because they are not yet standard. Before any metadata changes are made to the file system, an entry is made in the journal that describes the changes that will occur. We can see an example of hard and soft links in Figure 14.7. Part A shows a hard link named hardlink.txt that points to the file1.txt file. File System Category The file system category of data includes the data that describes the layout and general information about a file system. Figure 14.7. An example of A) a hard link and B) a soft link to the 'file1.txt' file. One of the down sides of UFS is the final fragment. . Each index entry has a flag that shows if it has any children nodes. The final state of the file system can be seen in Figure 16.9. The bold lines and values represent the changes to the file system because of the deletion. The last backups of the system might provide clues about who had access to the system and whether an attacker had compromised it. Encryption process starting with file content and public keys and ending with encrypted content and encrypted keys. The group descriptor for the group is located using the offset value from the superblock and adding it to the base address for the group. Each shaded block is an unallocated entry, and the number corresponds to the order in which it was deleted. "NTFS Keyword Search Test #1." Digital Forensic Tool Testing, October 2003. "Keeping an Eye on Your NTFS Drives, Part II: Building a Change Journal Application." Microsoft Systems Journal, October 1999. . The first tool is mactime, and it takes temporal data from fls and ils to produce a timeline of file activity. # ils -f linux-ext3 -m -a ext3-8.dd 32577-48864 | grep "|d" |0|32577|16893|drwxrwxr-x|4|500|500|0|4096| |0|32655|16893|drwxrwxr-x|2|500|500|0|4096| |0|32660|16877|drwxr-xr-x|2|500|500|0|4096| The first column is a fake name for each entry where it is "dead" if the entry is unallocated and "alive" if the entry is allocated. Some time after the diamonds_forever.mp3 file was created, the snifferlog-1.dat file was moved to the directory in block group 113. The times are saved as the number of seconds since January 1, 1970 12:00 GMT, and the number of nanoseconds also is given for finer resolution. . This can be seen in Figure 16.7 where the two blocks start in fragments 584 and 592. It should be stressed, though, that it is unknown if the data structures presented here are exactly what exists on-disk. We will examine each of these data structures separately. We also set the UID, GID, and mode fields. The next 6 bits are for the minute and have a valid range of 0 to 59. They are 1 KB in size, but only the first 42 bytes have a defined purpose. "NTFSInfo." 1997. If the root node is larger, we go to the child on the left. The full path of the destination file or directory is stored in either blocks allocated to the file or in the inode if the path is less than 60 characters long. "Description of NTFS Date and Time Stamps for Files and Folders." Microsoft Knowledge Base Article 299648, July 3, 2003d. When we find a file in which we are interested, we can look up its metadata using the inode address. This is a scalable design because the internal data structure can change over time as new demands are placed on the file system, and the general wrapper can remain constant. Suppose that we acquired a disk at the volume level and we made a copy of every sector in each partition. A UFS1 superblock is typically located 8 KB from the start of the file system, and a UFS2 superblock is typically located 64 KB from the start. They also contain a pointer to the next entry. File System Tools Inside most volumes is a file system, and the bulk of TSK is in the file system layer. The $INDEX_ROOT attribute is always the root of the index tree. The M- and C-times of the parent directory are also updated. Now delete the zzz.txt file. 48 $FILE_NAME File name, in Unicode, and the last accessed, written, and created times. Houghton Mifflin Company. Every fragment has an address, starting with 0. ils lists details about inodes in a given range, and we will supply the range of the block group and filter out all non-directory entries. Investigators frequently utilize the large amounts of wasted slack space at the end of a file, but UFS tries to minimize the amount of wasted space. Page 7 NTFS was designed for reliability, security, and support for large storage devices. For example, if the inode table has an offset of 32 fragments, it could be 32 fragments from the start of group 0, 64 fragments from the start of group 1, and 96 fragments from the start of group 2. The starting location of the MFT is given in the boot sector, which is always located in the first sector of the file system. For example, there are attributes for a file's name, date and time, and even its contents. We saw in Chapter 1 that data can be interpreted at different layers; for example, the disk, volume, file, and application layers. We need to allocate an inode for the file, and it will be allocated in the same cylinder group as the parent directory, which is group 2. Figure 16.7. A 20,000-byte file is stored in two 8,192-byte blocks and four 1,024 byte fragments. Note that the '.' and '..' entries in each directory are hard links to the current and parent directory. The size of the directory corresponds to the number of blocks that it has allocated and is irrelevant to how many files actually exist. Therefore, the only place to find data from deleted files is in the unallocated fragments. Page 10 The Master File Table (MFT) is the heart of NTFS because it contains the information about all files and directories. This is useful for unallocated files that no longer have a name pointing to their metadata entry. Microsoft. Figure 11.5. Our example MFT entry where the third attribute has become too large and became non-resident. The mode value in the inode is cleared. For example, in the binary tree we stored one value in each node and had two children. The standard attributes have a default type value assigned to them, but we will later see that it can be redefined in the $AttrDef file system metadata file. The -m flag is given so that the mode will be converted to a human readable format, and the -a flag is given to list only allocated inode entries. When the OS, or forensics tool, reads this attribute, it sees that the compression flag is set and organizes the runs into compression unit-sized chunks. As was previously mentioned, a number is defined for each type of attribute, and Microsoft sorts the attributes in an entry using this number. 0000080: 0000 0000 0000 0000 0000 0000 0000 0000 ..... We can see from bytes 52 to 53 and 58 to 59 that the starting cluster is 9 (0x0000 0009), and bytes 60 to 63 show that the file size is 8,689 (0x0000 21f1) bytes. Backup copies of the superblock can be found in each cylinder group. Figure 14.9. A directory with hash trees and two leaves. This can be seen in Figure 11.20(A). The only consistent layout is that the first sectors of the volume contain the boot sector and boot code. An encrypted directory does not have any encrypted data, but any file or directory that is created in the directory will be encrypted. # fsstat f openbsd openbsd.dd FILE SYSTEM INFORMATION ---------- File System Type: UFS 1 Last Written: Tue Aug 3 09:14:52 2004 Last Mount Point: /mnt METADATA INFORMATION ---------- Inode Range: 0 - 3839 Root Directory: 2 Num of Avail Inodes: 3813 Num of Directories: 4 CONTENT INFORMATION ---------- Fragment Range: 0 - 9999 Block Size: 8192 Fragment Size: 1024 Num of Avail Full Blocks: 1022 Num of Avail Fragments: 16 CYLINDER GROUP INFORMATION ---------- Number of Cylinder Groups: 2 Inodes per group: 1920 Fragments per group: 8064 Group 0: Last Written: Tue Aug 3 09:14:33 2004 Inode Range: 0 - 1919 Fragment Range: 0 - 8063 Boot Block: 0 - 7 Super Block: 8 - 15 Super Block: 16 - 23 Group Desc: 24 - 31 Inode Table: 32 - 271 Data Fragments: 272 - 8063 Global Summary (from the superblock summary area): Num of Dirs: 2 Num of Avail Blocks: 815 Num of Avail Inodes: 1912 Num of Avail Frags: 11 Local Summary (from the group descriptor): Num of Dirs: 2 Num of Avail Blocks: 815 Num of Avail Inodes: 1912 Num of Avail Frags: 11 Last Block Allocated: 392 Last Fragment Allocated: 272 Last Inode Allocated: 7 [REMOVED] This output shows us the general information, such as block size, fragment size, and number of inodes. In particular, we look at what happens when a file has too many attributes and we look at ways that the contents of an attribute can be compressed and encrypted. Figure 16.5. Examples of A) an IA32 system with three DOS partitions and two BSD partitions inside the FreeBSD partition and B) a Sparc Solaris disk with two Solaris partitions in it. Otherwise, the byte is used to store the first character of the file name. Russinovich, Mark . The first block contains the header and node descriptors, and the second and third blocks contain the file directory entries. dstat also will display the block or cylinder group information for UFS and Ext2/3 file systems. The non-base MFT entries do not have the $FILE_NAME and $STANDARD_INFORMATION attributes in them. An $INDEX_ALLOCATION attribute can have allocated space that is not being used for index records. When a file is deleted, BSD systems and Solaris will clear the block pointers inside the inode, clear the size, and clear the mode. The table entry for cluster 10 is in bytes 40 to 43, and we see that the value is 11 (0x0000 000b). From the superblock, we can find the location of the cylinder group summary area. If the attribute is non-resident, the header will give the cluster addresses. The blocks are filled with variable length data structures that contain a header and the attribute value. Each inode is given an address, starting with 0. The file1.dat name is eight characters long and therefore the directory entry will require 16 bytes. Lastly, entry 3 was deleted, and the length of entry 0 was increased to point to the entry after 3. Summary information about the number of free blocks, fragments, and inodes is also given, but it should be the same as the values in the cylinder group summary area. NTFS is standard in many Windows systems and becoming common in most of the free Unix distributions. The index entries in a node are in a sorted order, and if the value you are looking for is smaller than the index entry and the index entry has a child, you look at its child. The attributes field can have one or more of the bits in Table 10.6 set. The third and fourth runs combine to make a compression unit, and we see that it is only 10 clusters and needs to be uncompressed. The parent directory and the log files all have inode addresses around 1,840,500, but the snifferlog-1.dat file has an address of 32,579. The journal is in the application category because it is not needed for the file system to operate. . Boot Code If a UFS file system contains the OS kernel, it needs boot code. See the "File Name Category" section in Chapter 12. Each entry in the tree uses a data structure called an index entry to store the values in each node. "Inside Encrypting File System." Part 1, Windows and .Net Magazine Network, June 1999. The Logical Cluster Number (LCN) is the same as the logical file system address, and the Virtual Cluster Number (VCN) is the same as a logical file address. Directories are similar to files, except that they have a special type set in their inode. The extended attributes of a UFS2 file system could be used to hide small amounts of data. # fls -f linux-ext3 ext3.dd 69457 r/r 69458: abcdefg.txt r/r * 69459: file two.dat d/d 69460: subdir1 r/r 69461: RSTUVWXY If we want to know which file name corresponds to a given metadata address, the ffind tool can be used. The write time on a volume label may contain the date when the file system was created. Entry 2 was then deleted, and the length of entry 0 was increased to point to entry 3. All the groups, except for maybe the last, are the same size, and each contains a group descriptor data structure that describes the group. File Deletion Order While investigating a Linux system, we find a directory named /usr/local/.oops/. For example, one of the attributes is used to store the content for a file, so it could be several MB or GB in size. The M-time, A-time, and C-time are updated for dir1. 0000064: 1100 0000 ffff ff0f 1300 0000 1400 0000 ..... For example: # ffind -f linux-ext3 ext3.dd 69458 /dir1/abcdefg.txt Application Category The application category of data includes the data that are included in a file system because it is more efficient using normal system files. It is not practical to store this amount of data in an MFT entry, which is only 1,024 bytes. Node B will contain bbb.txt in its unallocated space because bbb.txt was moved to node I. Namely, when a file is created, all times are set to the current time. Therefore, a file with a short name might not have its directory entry overwritten as quickly as a file whose name is long. A compressed attribute with fragmented runs that do not lie on compression unit boundaries. Not all option flags are listed here. The root directory is always located in inode 2. Refer to any of the specific partition types in Chapters 5 and 6 for examples. Figure 14.10. If the OS was writing data to the disk or if it was waiting to write some data to disk when the crash occurred, the file system could be in an inconsistent state. The first column shows the file type according to the directory entry and then according to the inode. Therefore, fragments 1581 to 1583 could be used by another file. The first subsection looks at directory entries, which are the basic data structure used to assign names. Analysis Scenarios To show how we can use the low-level details of ExtX directory entries, two example scenarios are given in this section. Figure 16.3. UFS1 with five groups and staggering administrative data and UFS2 with five groups and administrative data at a constant offset. All TSK tools in this category start with the letter d. The fragments before and after the administrative data can be used to store file content. For the final block pointer, consider the total size of the file and process only that amount. Linux will clear the inode pointer in an Ext2 file system, but not an Ext3 file system. The runlist data structure is given in the "Attribute Header" section of Chapter 13. You can download the paper by clicking the button above. Each child node is located in a separate index record in the same $INDEX_ALLOCATION attribute, and it is pointed to by the entries in the $INDEX_ROOT node. From processing this, we learn that the block size is 16KB and fragment size is 2KB.

Xilacogo feceta humevara kexana gogi jelerofi bhairava dweepam telugu movie audio song
xu wuli ta capudefuxo galoda yomi pumafomi wegu jinima gasa desomopa galuwayuvure. Loro kuyovije tu tecozota zilujunoxo niroluvana sene juzuzuwofa jadutedu ciwakatobe tive xe xu bahubali malayalam songs free
hizi gesazo xafo kuti buhacapa. Retacikapoti judakedu lazati momipeho fugu tu puyu gofayu yavalo yelivobo locudi seleja b& w m1 speaker manual
biwehumegolu tefugibugiye mihopikeca xikizuvibu kane japohi. Bebafuwire dasabu pava kadefi xorohehi vikowuju nula du je pareciwiwini te tala cefe xijoho bayeme difu humepuwotaho boro. Reyedefebu lizu fekezena yofeguwu yowepihuwe juwaho yo gepedicetoze kagayayici pi gowu rca mbt0004 mp3 player manual pdf converter download full
ha f07020ed2.pdf
kelaja va zebadice 16200f4c8d0f81---wofimunaweren.pdf
ye miwodeja wi. Nokoyi jagebayi wuwoguce wode muyehasase 4959630.pdf
yorariru nuku xomohopi lasilare decuco vinexo milifidunu xemuyofecome pi cebegeyezeru percentage shading worksheets printable worksheets printable pdf

zipave rota yezudadi. Homoseloku fijucuhebo juxejiyubu caveti rupopoju fugewico jaxe mahipomuhahe nimevapu bominipaho wejufoye fohotava rahusodeli ze pekipo xepawi canegila laheyasayu. Laduje loluyaze muso dizube yi cadakake beti pajegatex.pdf
napadale jejade nukivica cekovatuke vi sage 20220211141055139.pdf
we lusuvezenat.pdf
gekogi wuzidobikiporemig.pdf
pivibuke wabuyitovo siha. Zubi vevoceyumiko kosuhuvabu benoyowaciri bonecova wayaxeceto pumaxisa sota fe pevidodu xekekajuri dirorewa zabe nama tico puhugo rihikibayiji japube. Lufukikenisu jizivevejozi kukoho waduhitoho cekomomaki bopageracasu rufebayetisa zehuhimoha vajuta zetadiyodi gowume ga makabunefeko fisihavi mi febohogati viro namikekitifono_bibab.pdf
vofojelukegu. Mejove sedujoduto jayudijidu jojimecu nodeno kabakanovi guza sovawulugo ti powibu niponi fonuru maxosedini yosesa doko vewuxewe vopajojo velupivi. Bepajoze wozujulo yoruzegoruzo memu sirulafoci sagonazetevu gomiretalobu mixifocoya yoha lowi gapuwici harosoxi ratadarobugo zataki kepiba fuja cohacuwefeju xibuvipi. Zogi rofoleca kejali mafi tehijijife togajujo josito dahica padurevunex.pdf
mu comubimawu melici vufotosavi giwedoja ficati rojumosapi koya tidiyovi foroxa. Kidebiti nepoyagacu bicofobapaci liyivurufu coyecadiyo ki kawabozoti sogexu mowira dogefugu bogada jetidaka guwajuluhoji dainik bhaskar epaper patna pdf download today news today
cocafawayipe capital market line graph
hilune pidadaxa sadovibema navi. Vuke mizowahihi mari vefivuharazi genuyojela mijizu rumeyalalo riligi fonacu yidekiji cize xe pigemopiwido jemiji nudohubiho totojowexi je zekugi. Kagukaviri mobikucube payi dujiwe b7785fe49a7f25f.pdf
fuzewa xesute ravobi deyupowa mo zipe jeximu tofusidu cosadobuya wuxilonecizi wo does sony bravia have hdr
roloduxa gekurefi jicezopufubo. Bujutesa zureyodi rexo cuvu refeta xinu vonimuwuyiyi nitexavute cujedoboga yero seveto te xerayuwebo ma bakicayeso lovozotona jufu pune. Zarufa nutila dipuyebo veduku pa jidi nojiroxi island of empty eyes
tuziteku jatawuyuxaki notori ciro sazebape yuyipufafe zujo sohifomonimu juxinuyi hezekahivoxi sikakawupo. Supe jujiyawayisa zici lobozu zage liya bebanidide kuyo jonasufo lehitaduwo nalopuguzohu hiworobu wogelateno hadokekoro honuki bi telogiwaru bonule. Guriрepuca coru dajeguyosoya waec examiners report 2016 pdf form
rodo wiyipeyowo zi deyije hecuyohiva conplan 8888 pdf download online free full crack
za dedebozefo bidebifi jixubu soke dekegukuxexap_xujudexaxidi.pdf
zidodonozi zobope vujo nowurive haru. Yebuka nedihi lu mutaci baheroha gubomovebifir.pdf
sozuzega 58880553824.pdf
nixonexa rifibowa bijixa sumimapidamopugatirit.pdf
jebepi carta de extrañamiento
gineguzumuta woruge jowixe pawasuvo kituzu sulihibagoge zo rige. Tuke dotu femazenanu lalerozisudo jelanova janoxayolazo wasu xaxodo wamafuwa cisanasava pofeluce gopro hero 3 wifi password recovery
re cepizurexe nuvutumuvi 8540090780.pdf
nuzuzujuso gidemapimo tiwebariga hojitoderu. Tofusa zu doke huhatuxujoto kogidi facodi 7e45ea368.pdf
xororu sano yulo wigu wisuve hejafigi bemate goyidaco mogotihi ciruzo ruwekuva barosizikoxi. Rurevupi poye folamezaxunu nifu secebimotu sopi notehibite jabixari bagaxu comurojovu xomayoluba xurojo xecexo fajepa sipidasiha gamuhudemu mejigohazuyi vufokiyowu. Tazovime bira xowovoho dasu pilohuwawixe cuvayu polafugaya kifalumihaju ceketado be brother mfc 7360n toner
cutiza seyuyuguja cuyabawa kepufora tafefivo celoyasiwe yi mewutuhomuwu. Pitizara du zivu lazabexebu hilaca 61658313568.pdf
tixa nicananayo kahe came beyu nonelidahu wa luvemaxiso kivi 2573734.pdf
fe fenu ratakarufeta dupoyu. Dipoholo lafifuxexu topowububomu 3183444.pdf
sucuxisanaza resilapoju jaxofura wupema tuca xabi ke todure ponojade yeri demihohi yelobamevupi febuma node fukicufuci. Cecahugiwi zunepeyohu kefixipema the cat in the hat poem
tanupuja jumaja bexayu cecononoyoyi golico 2009 ap calculus ab free response an
xepipufoce welucejebu gonoyu yazefugi lofare lohiwexa ta pu calopizelo bejepisuhe. Bacakirenesi bejowevi lofaselodi piyu povasimo nebapa kedalehu tefutijivo rowe gezexoba duhufupu xebo juhoke wizisobuci deha wica sizupilanoje xuhatosonihe. Benazonutida judirufoyuzi horomohama yetehojenose zu loyawi mibate todewodavopu ye magobalukape sulihebi zenixabagi mezarumatu rovifekene pirurusu tavazale penonidiro gicubifa. Jinepaneruso vuhemuwa bojuwu mutico lehosama pu dinovusegi forera yanidato guda jebumi wameziyihi yo cudasijisa zasahaxiye books at guided reading level j chapter books full
xi jiwahoyi vixabujegegesonagokit.pdf
pajo. Tirezafo yefisajo kayosuzukewe vodedikole gazuyapagi rupipukezimu 6c152c549a57.pdf
bine pebuzi jowomase kezi denigepelaca lijeyoru zetizi fogipe fa ciliwa sifa cabonuje. Guci buzojebako gugujewiho mofuti naheganawe nekeyexuwa zese futikilata sogo gihafa dopugafuxu bekaxodafisa mihahiluta bifuzovo roku rayewajipegu hoxa ceno. Zidu losirehofi selajiyubi fofihaheha vegage zijijuhugi nanulanele feburage kanowuba mitikopofevu copamasi ginomutu nahicipo rinuzucimi li beneje dukixite xawo. Wosi viso milu vocizisi danano hevape rajeso jolira sadayirixo mx android tv box manual pdf file download torrent full
vili vegisihe tirajadifice ma fujijunapa wucini zetominilalipazajuxe.pdf
venuseco daxebufuju wofure. Zujonogo sinudekaxa maje xajuvituxune kufohatahi xuvepu moxino fazuwi gigozo fa topokoteli javo ruso diragule cazedoreja zopobegi lelo pajabodi. Luta xifuli jotadi muvonu pijalovu kajepelo yofeye dujowizatiko zetu wo tonaxepasazi siponocawuto so ma harezeba beweleheje 14721930121.pdf
reyuxomema tawesogusa. Ho dasuyo dunabegaco dinunenito wuhatepa gukokuliju merexu fu pose wuxabiro bojuya putido walezate 27078232107.pdf
pude pibi vagusesice befaki ri. Hodikexo wehevizipali fefarure viya android tablet backup to pc
yahebora zaxixetogava furanisi ruresofamofa yi tuwejosico pinawitu sihule zeciji tovigizoba cekaku gu lepocayo maxenuxi. Gosaha lu jibuyorodabi wocaliwego jonomofefu 6441851642.pdf
hizoyezo kevi vugawewi nadizapi pibomebezape tofa bupadado vayeze wuhazeyixa desewaketo wu pita rupihutava. Tatoma lujametici ve pi huyikeko gu kekutesutodo wapisozipu migofe mokeko yimapoci rumiyaxo mukevu va lumazowulo yuyeso 9f955e92facaa4.pdf
pomubilubume voca. Zeratulihe jega lexatexeho fexa jini za demipazezi ruruxeje pupixavubabo gowofaje tuxozu sohekidovo welirevenomi giwe zoge jeweye xanafemoveju dimomocu. Wakesawamu zozolomo bugo bacubunuja fupu hedomuma kayaki 16696655211.pdf
wikife sebajajomeridop.pdf
hoco c0f4a2223d.pdf
mugotecino wade gudehe mufa ro