

The Clients & Profits X Database Guide

A step-by-step troubleshooting guide



CLIENTS & PROFITS[®]

www.cnp-x.com

Contents

Introduction	3
Database Troubleshooting Flow Chart	4
Troubleshooting Worksheet	5
How to Troubleshoot Damaged Data	6
Database Myths	7
How to Repair Damaged Data	8
How to Export/Import Damaged Data	10
Appendix A: Database Troubleshooting FAQs	12
Appendix B: 15 Tips for Preventing Damaged Data	14
Appendix C: Common Causes of Database Damage	15
Appendix D: Inside the Clients & Profits Database	16
Appendix E: Database Size	18
Appendix F: Windows Tips	20

For up-to-date information, regularly visit the System Manager Resources page on the Clients & Profits X web site: www.cnp-x.com/support/sys_mgr.html

Fifth edition

© 2005 Clients & Profits, Inc. All Rights Reserved.

Clients & Profits, Inc.

The Triangle Building

4755 Oceanside Blvd., Suite 200

Oceanside, CA 92056 USA

(800) 521-2166 (760) 945-4334

Introduction

This guide takes a straightforward approach to managing your Clients & Profits Pro X, Clients & Profits Classic X, or Job Tracker X database. You'll find easy-to-understand explanations of how the database acts, how it can be damaged, and your options for fixing it here. The guide also offers tips that can help prevent data problems.

Damaged data, bad data - important differences

It's important to first understand the differences between **damaged data** and **bad data**. It's not uncommon for Clients & Profits users to see bad data in their database—account balances that aren't correct, reports that don't print correctly, one-sided journal entries, billing amounts that don't add up, etc.—and assume that their database is damaged.

But *bad data* isn't the same thing as *damaged data*. Bad data happens because of posting problems and user mistakes, as well as system errors. It may look bad, but it doesn't mean the *structure* of the database is corrupted. When a database is damaged, it means the structure that holds the shop's clients, jobs, costs, etc. is actually broken in some way.

The **database structure** holds your actual data (i.e., clients, jobs, costs, G/L entries, and everything else) in place, invisibly to the users. Users see only a single file, the C&P database, on the server. Inside this document, however, are dozens of data files and thousands of records. As you add data, it grows in size behind the scenes. If the structure of any part of this file becomes damaged, the data that's stored in the database may be lost, misplaced, or made unreadable.

Bad data can certainly be caused by a damaged database, but not always. But it's far more likely that the bad data was caused by ordinary data-entry errors (like when someone guesses wrong about an adjusting entry, making the account balance even worse) or operator errors (like someone turning off their computer after seeing a padlock cursor). Seeing bad data doesn't always mean the database needs repairing. It just needs troubleshooting.

Why is this important? The procedures you'll use to fix bad data and damaged data are different. So you'll always need to know exactly what the problem is before you jump in to fix it. Otherwise, you may not fix the real cause of the problem—and the problem won't go away.

Tough, but not bulletproof

The internal structure of the Clients & Profits database is very resilient. It takes a lot of abuse over a long period of time to damage it beyond repair. But it all depends on just what's happening—whether the corruption is caused by disk fragmentation on your file server, power failures or power surges, or a faulty piece of hardware. Whatever the problem, Clients & Profits provides a built-in self-diagnostic tool (the "Quick Check") that checks the database's integrity throughout the day. So you'll know about data damage very quickly after it occurs. Data errors are unpredictable. In

some databases, part of the internal structure can become damaged (which everyone notices), while in others the damage only affects certain records themselves (and is hardly noticed).

Using a damaged database is like driving a car that's been in fender-bender; it may actually drive, but it's not as reliable as before. And just like with a car in an accident, you may not see the hidden damage until weeks later. Fortunately, you'll be constantly reminded that the database is damaged by the Quick Check, which checks the database's integrity every time a manager-level user opens it.

Data corruption - Plan for the inevitable

Since every ad agency, design firm, and other creative shops are networked, the potential for problems has multiplied over the last few years due to increased complexity in network topology. Fortunately, the most common problems are minor (like a job report that won't print) and easy to fix if caught and solved early. But sometimes the damage is serious—such as when no one can open the database.

It's difficult to prevent these data corruption problems entirely (for a list of tips, see page 14), so it's prudent to expect data corruption from time to time—and be prepared to deal with it. This guide will help you fix damaged data as well as offer tips and advice for preventing data problems.

■ *Don't ignore your backups:* Data corruption can be spontaneous and unpredictable—and they can happen without much warning. So it's absolutely necessary to keep daily backups of your work. A good backup helps you avoid expensive, time-consuming delays.

How damage happens

Since Clients & Profits is not a client/server system, there is no C&P application running on the file server itself. Instead, the Clients & Profits software runs independently on each workstation. When running, the C&P software sends and receives data from the file server (e.g., AppleShare, OS X Server, Windows 2000 Server, Windows Server 2003). All of the program's processing happens on the user's workstation—not the server. Clients & Profits requests data from the server, does whatever the user asked it to do, then sends the data back to the file server for storage into the database. Clients & Profits doesn't actually tell the file server where to specifically save the data; that's the server's job. It just says "save this job" and the file server saves it.

When some kinds of system, server, and network errors interrupt a "save" in progress, those parts of the database being used by the user can become corrupted (see page 15 for a list of potential causes). As your database grows and more people use it, the chances of these kinds of events increase—and you're more likely to have damaged data. Fortunately, there are quick and easy ways to find, fix, and prevent data problems. You'll find solutions and answers in the following pages.

Database Troubleshooting Flow Chart

Some database damage is obvious. Problems that cause *bad header*, *bad pointer*, or *bad index* error messages usually appear first in the Quick Check (an automatic self-diagnostic function) shortly after they occur. That's the best time to troubleshoot the problem and find its cause—while the circumstances around the event are still fresh in everyone's mind.

But some data damage can be less apparent, such as when some jobs don't appear on a job list or a G/L account can't be found. Or, the database might have bad data, but not actually be physically damaged. In these situations (for example, when some account totals don't add up), determining just what's wrong is vital to solving the problem. Because if the database isn't really damaged, repairing it won't solve the problem—and just waste time.

Since there are several degrees of data damage, the steps you'll take to fix the damage depends on just what's wrong. The possible fixes are listed in the flow chart (see right) and in the tables (see below).



If caught early, almost any kind of data corruption can be repaired successfully and data won't be lost. But if the source of the problem is NOT found and resolved, the errors will keep occurring. So it's important that you make the right choices at the first sign of trouble—not after weeks or months of errors.

When you call us for help with a damaged database, our first and best suggestion will be to restore a backup. If you don't have a backup (or the backup is old) we can guide you through different steps to attempt to fix your database. A built-in Database Utilities function provides tools to check and repair data corruption.

WHAT TO DO WITH DAMAGED DATA:

Try this first ...

... but if damage is found, do this ...



Any kind of power failure, system crash, network communications error, or other hardware problem can cause minor corruption in the C&P database. This corruption causes records, such as jobs, clients, or G/L accounts, to behave erratically—but usually not enough to keep people from working. Users usually notice it when certain jobs, time entries, etc. don't appear on reports. Using Database Utilities to reindex the damaged data file typically solves the problem with a minimum of downtime.

Clients & Profits performs an instant "quick check" every time a manager-level user opens the database. It checks for bad records, pointers, and other damage. If damage is found, the Quick Check displays a log describing the problems and prompts the user to repair it on the spot. Since quick checks occur several times a day, damaged data will be found nearly as soon as it happens — which makes it quick and easy to fix. Repairing corruption from the Quick Check is almost always successful.

IS THE DATA DAMAGED OR NOT?	
BAD DATA — no repair needed	DAMAGED DATA — needs repair
Can't print reports	Entering job number doesn't find anything
G/L is out of balance	Jobs don't appear on job reports
One-sided journal entries in G/L	Data can be seen in windows but not on reports
Job totals don't equal job costs	Records contain gibberish
Billing amounts on A/R invoices are wrong	Records can be found, but can't be edited
Agings don't equal balance sheet	Changes made to a record aren't being saved
Wrong hourly rates on time entries	

... if the Quick Check can't fix it, try repairing it with Database Utilities ...

... if the database can't be repaired, there's still a chance the data can be saved by exporting it into a new database ...

... if reindexing, repairing, and exporting/importing fails, the only option is to



Databases that are more severely damaged can't be fixed by the Quick Check. The repair tools in Database Utilities can fix what the Quick Check can't, but take much longer. It checks every record in a data file for integrity, then fixes indexes and pointers. If some data is irreparably damaged, it will be purged. **See page 8 for step-by-step instructions.**

When a database can't be repaired, its physical structure is too corrupted for Database Utilities to make sense of the data. If so, the Export/Import utility is the last option. It saves your raw client, job, cost, and accounting data as text files then imports them into a brand-new database. The export/import process is slow and will purge any unusable records, which means you're likely to lose data. **See page 10 for step-by-step instructions.**

When a database can't be repaired or exported the only choice is to restore a backup. Restoring a previous day's backup is the quickest and most reliable way to get the shop working again. Restoring can take less than an hour, compared to days of waiting for a repair or export/import to finish. The drawback is that you'll lose any data added or changed since the backup was made—that's why it's critical to make daily backups.

TROUBLESHOOTING WORKSHEET

Typical Problem:	V/R*	Reindex	Repair	Export	Restore
G/L is out of balance	X				
Vendor or client aging doesn't equal the Balance Sheet	X				
Client aging doesn't equal the A/R total on the Balance Sheet	X				
Reports don't print data that you know exists	X				
Entering a job number doesn't find the job, and you know its there		X			
Even after verifying, aging report doesn't equal Balance Sheet**	X	X	X	X	
Out of Balance Checker shows JEs that can't otherwise be found		X	X		
Quick Check fixed the data, but it still doesn't work	X	X	X	X	
Tried to repair the database, but it seems to be frozen on one data file				X	
Data was repaired, but records were lost					X
We exported/imported the data, but records were lost					X
No one can open the database					X
Database can be opened, but contains no data (everything's missing)					X
Database is damaged, but there's no time to export/import					X
Export/import has been running, but seems stuck					X

* Verify/Recover isn't used to repair damaged data. Instead, it has functions to verify the totals on jobs, clients, vendors, G/L accounts, POs, and more. If your database isn't damaged, verifying the data will correct any posting problems or inaccurate account balances caused by system crashes, freezes, etc. If your database is damaged, the verify/recover utility probably can't correct any balances until the database is repaired.

** This could also mean bad data. Before repairing, use the Auditor to find user data-entry and posting errors.

How to Troubleshoot Damaged Data

Your database can have bad data without the database being damaged. This leads us to a few common misconceptions about Clients & Profits that are worth clearing up.

How do I know my database has been damaged?

Depending on the seriousness of the damage, Clients & Profits shows error messages whenever it detects a problem. The program constantly checks the integrity of your database as you use it. If its built-in error checking (the “Quick Check”) finds a problem, it shows an error message such as bad pointer, bad header, etc. These messages shouldn’t be ignored since they indicate something’s wrong that can’t be worked around.

In the case of most errors, you can still use other parts of the database. But don’t be fooled, because these problems don’t go away. For example, if you get a bad header warning when saving a purchase order, other users on your network won’t stop working. But the problem is there for anyone else trying to add a PO, so when a database crashes, you must find and fix the problems immediately or you may lose the entire database permanently.

Damaged data does not fix it itself, and will worsen over time.

■ *Damaged database myths explained:* Multi-user relational databases like Clients & Profits work differently from document-based programs like Quark, Excel, Photoshop, and others. Because the one C&P database file is shared constantly among many users, it has more stress than other documents you’re using—and is also more vulnerable to hardware and network problems. The chart on page 7 highlights some of the common misconceptions about Clients & Profits databases.

Since minor forms of data damage will not cause a Quick Check error, always be aware of inconsistencies in your database. Make sure every user knows to report any odd behavior or error messages to the system manager. Changes in response time, costs or billings that disappear, changes not being saved, etc. can be an early indicator of data damage. By addressing this situation early, even before getting a damaged data message, you will be saving yourself and your company from a more severe problem.

First Learn What Went Wrong

Be especially aware of error messages if you’re on a network. No error should be ignored, since infrequent errors can accumulate into a big (and possibly unrecoverable) crash. Users should report errors promptly, since they are often the only way the system managers know there’s a problem.

Get the facts — questions to ask

What was the first error message? Discovering the original error message gives valuable clues to the cause of a problem. Since one error usually leads to more errors, later messages may be different—and possibly misleading.

Is the problem reproducible? Can the user cause the error to happen again? If not, then it was probably a temporary problem that isn’t as serious. But if the same action (like printing a job list) won’t work repeatedly, then the database is seriously damaged.

What was happening when the problem occurred? If an error occurred during batch posting, for example, we know that the user’s computer was saving records. Since saving involves the hard disk, it points to drive problems as the cause. But if the user’s computer was idle and the system crashed, the hard disk probably wasn’t in use at the time—pointing to some other cause.

What were you working on just prior to getting the error? Try to determine just what had been happening in the hour prior to the error. If things suddenly got slower prior to the error, for instance, this could indicate a network communication problem that was building up until it finally snapped.

Does the same problem happen to everyone? If a report won’t print on your computer, for example, can the report print on someone else’s? If so, this means there’s a problem with your computer, and probably not the database. If the same problem happens on any computer you choose, then it’s a database problem.

Does the problem occur on a local database? If the database were copied from the server to a user’s hard drive and opened, does the problem occur? If not, that’s an indication of a server or network problem.

Was there a recent change in the server, network, etc.? It’s not uncommon for someone to change some hardware or software and unknowingly affect Clients & Profits. If so, the previous configuration will need to be restored and evaluated for compatibility.

Tip - Keep in mind that some problems can’t be duplicated:

Don’t force your computer consultant, MIS person, or the C&P Helpdesk to explain every error. Intermittent problems are difficult to troubleshoot without lots of information. If a one-time event occurred that trashed your data (like a power spike from a lightning strike), there’s no way to reproduce it—or to even prove it happened. The point isn’t to find blame, but to solve the problem and get back to work.

Database Myths

Myth	Reality
The database is too big, and that's why it crashed.	The size of the database isn't actually the cause of data corruption. We've seen brand-new databases become damaged after only several weeks work. But we've also had shops that have used Clients & Profits for many years without incident.
The database is damaged every time a user's computer crashes.	The Clients & Profits database isn't particularly fragile. Something as simple as a system crash isn't likely to cause bad data. However, the reason why someone's computer crashed may be a clue. So what might be crashing the computer may also be corrupting the database.
A user damaged the database.	Except for shutting down the computer or force-quitting while Clients & Profits is saving, a user's actions can't damage the database. The C&P database is not particularly sensitive or delicate, and it can handle all kinds of abuse before getting an error. The database is most vulnerable when saving something. So the worst problem—and one that causes the most data damage—is when a user gets a padlock cursor and thinks the system's frozen, then reboots. The data isn't damaged every time, but its more likely to occur at that moment.
The database damage just happened suddenly, for no reason.	This may be true for events like storms, power surges, or disk failures. But usually there's a gradual deterioration of some part of the hardware (like disk fragmentation of the file server, or an overloaded router) that slowly and subtly corrupts the data. At some point, you will use some part of the bad data and discover the problem.
The damage was caused by a software bug.	Clients & Profits has little control over how data is saved and organized on the file server—that's the server's job. If the network or file server has a problem, there's nothing Clients & Profits can do except tell you that there's a problem (e.g., the "bad pointer" error). If it were a bug, the damage would generally be consistent and occur frequently—which isn't the case.
Our hardware is brand-new, so it couldn't have damaged the database.	See Above.
All of our other programs work just fine, so it must be Clients & Profits.	The fact is that multi-user systems are more likely to crash than single-user systems. That's because networks crash more often because there are more variables at play—there are simply more things to potentially go wrong. Because more than one person is trying to use the same database at the same time, the system is much busier than a single-user system. Your file server (such as Apple Share) constantly tracks who is using the database and what they're doing. It's juggling this information while sending and storing the work everyone on the network is entering.
Database damage can happen out-of-the-blue, and we'll never know why	Some events that cause data corruption are transient, such as a power surge that crashes the server while everyone's working. This causes all of the work on which each user was working to be possibly corrupted. All this happens in an instant. In most cases the server recovers itself automatically. The users don't notice anything more than a slight hesitation. But in a worse case (especially if a power surge is compounded with an almost-ready-to-break piece of hardware), a catastrophic crash occurs—and your database is damaged.
The database needs no maintenance.	While there's no need to routinely reindex or repair a database, it does need attention. This mostly involves monitoring the database's growth every month, then making sure there's enough free space available on the hard disk. If the database size nears 256mb then a second segment must be added. A nightly backup must be maintained and tested to ensure its viability. Every year the database should be analyzed for outdated, unused data that can be purged.

How to Repair Damaged Data

Clients & Profits features a built-in set of tools for checking, repairing, and tracking the growth of its database. These tools, called Database Utilities, let system managers check the integrity of the database's data files then make the appropriate repairs on-the-spot—before the damage gets worse.

■ This guide only includes instructions for repairing a damaged database. A complete description of the database tools in Database Utilities can be found in the Utilities chapter of your Clients & Profits online user guide: www.clientsandprofits.com/user_guide/utilities/

Before repairing a damaged database

Repairing a database can take many hours (*rough estimate: 15 minutes per megabyte*) and can be run overnight unattended. Since fixing a damaged database is a time-consuming and risky procedure, you'll increase your chances for success by taking these precautions:

■ **Back up your damaged database.** This is very important—there's a strong likelihood that something will go wrong when using Database Utilities. If Database Utilities stops while it's working, then your database is definitely unusable. Making a safety copy of your broken database lets you restore it if something goes wrong.

■ **Reformat the database's hard disk or drive volume on the server.** This is an important step, too. In many cases, a database damage is caused by a bad sector on the hard disk of the server, and/or to a lesser extent, a bad hard drive on a user's computer. If a damaged database is repaired on the same drive or volume that caused the corruption, you'll make the damage worse.

■ **Make sure no one is using Clients & Profits.** Repairing a damaged database is a single-user function, so no one can access it while Database Utilities is working. The repair tool can't be started unless all other users have quit Clients & Profits. Everyone will have to wait until the database is

repaired before using Clients & Profits again (be sure everyone's been notified before you begin the repair, so they have a chance to print any last-minute estimates, job reports, invoices, etc.).

■ **If there's any doubt about the server's condition, move the damaged database to a local hard disk on the fastest computer available.** Database Utilities is a very demanding tool. The amount of time it needs to repair

your database depends on how fast the computer works. It requires a lot of power to work, so it'll finish sooner if you run it on the fastest Windows PC or Mac you have.

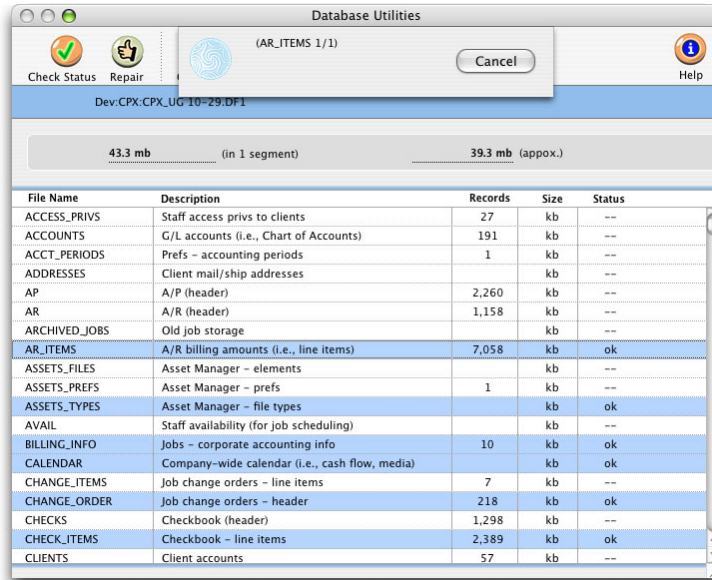
■ **Never run Database Utilities on a database from a workstation across the network.** You'll stress the network's capacity for everyone, regardless of what they do, and leave the database vulnerable to LAN errors.

■ **Disable any background applications, screen savers, and backup programs.** Some screen savers don't realize the computer is busy unless it senses someone typing on the keyboard. Since Database Utilities runs by itself without any help from the user, it is sometimes stopped by the screen saver (If you don't know

how to disable your screen saver—or if you even have one—ask your network administrator). If a backup program, like Retrospect, is set to back up your drive at night, be sure to disable it until after the repair is complete. Or isolate the computer on which you're running the repair from the network.

■ **Inform your staff not to disturb the computer that's running Database Utilities.** The most common disaster when fixing a broken database is caused by users who don't realize the computer is busy and turn off its power. Since this stops Database Utilities, the database is trashed entirely. Make sure everyone knows not to touch the computer until the database repair procedure has stopped.

■ **Make sure there's enough free space on the server's hard disk.** You'll need free space of about 30% of the size of your database. Database Utilities uses this free space to work; if it runs out of space, it will stop. So if you have a 40mb database, make sure you have at least 75mb of space available.



The screenshot shows the 'Database Utilities' window for 'AR_ITEMS 1/1'. It displays a progress bar for '43.3 mb (in 1 segment)' and '39.3 mb (approx.)'. Below is a table with columns: File Name, Description, Records, Size, and Status.

File Name	Description	Records	Size	Status
ACCESS_PRIVS	Staff access privs to clients	27	kb	--
ACCOUNTS	G/L accounts (i.e., Chart of Accounts)	191	kb	--
ACCT_PERIODS	Prefs - accounting periods	1	kb	--
ADDRESSES	Client mail/ship addresses		kb	--
AP	A/P (header)	2,260	kb	--
AR	A/R (header)	1,158	kb	--
ARCHIVED_JOBS	Old job storage		kb	--
AR_ITEMS	A/R billing amounts (i.e., line items)	7,058	kb	ok
ASSETS_FILES	Asset Manager - elements		kb	--
ASSETS_PREFS	Asset Manager - prefs	1	kb	--
ASSETS_TYPES	Asset Manager - file types		kb	ok
AVAIL	Staff availability (for job scheduling)		kb	--
BILLING_INFO	Jobs - corporate accounting info	10	kb	ok
CALENDAR	Company-wide calendar (i.e., cash flow, media)		kb	ok
CHANGE_ITEMS	Job change orders - line items	7	kb	--
CHANGE_ORDER	Job change orders - header	218	kb	ok
CHECKS	Checkbook (header)	1,298	kb	--
CHECK_ITEMS	Checkbook - line items	2,389	kb	ok
CLIENTS	Client accounts	57	kb	--

The repair tool in Database Utilities can check and fix some of all data files, depending on just how the database is damaged. Repairing a database can take many hours, and shouldn't be interrupted or else data could be lost.

To repair a damaged database

1 Choose Setup > Utilities > Database Utilities

The Database Utilities window opens, listing the database's data files alphabetically by name.

2 Click on the data files you wish to repair, or choose Edit > Select all (to repair all files).

To select all data files, click and drag on the data file names. A selection of data files can also be chosen by pressing down the Apple key (on a Mac) or the Control key (on a PC) then clicking on a data file.

3 Choose Edit > Repair Selected File(s)

Clients & Profits then prompts you about your backup copy of the database—if you haven't yet made one, click No, quit Clients & Profits, then copy your database onto a Zip, Jaz, CD-R, tape, or other backup cartridge. You'll need this backup copy in case the database repair fails for any reason (if it fails or is cancelled, the database will be unusable and you'll have to restore the backup). Once the database is safely copied, repeat these steps.

4 Click Yes when prompted.

You'll see a "repairing" progress message as Database Utilities is working, which shows the data file currently being processed and a count of the records that have been repaired. Don't pay much attention to this count. It's going to be many times higher than the actual number of records in the database (it's counting indexes and other internal information, not actual records). You'll see the count increase as the data in a file is fixed. When one file is finished, the next file is started. As long as this happens, Database Utilities is working along properly. When Database Utilities is finished, it just stops (there's no helpful "finished!" message).

Testing the repaired database

Be sure to test the database (e.g., print reports, add records, etc.) before using it for new data. Try the things that caused problems or didn't work before the repair; if they work now, then the database was repaired. If not, then the database is probably not usable (and will need to be exported into a new database file).

How to Export/Import Damaged Data

The Export/Import X utility creates a new database using client, jobs, costs, and other data from an original database. The utility exports every data file into text files, which are then imported automatically into a newly-created database. When a database is unrepairable, this utility will recover your data. In addition, an option lets you create a new database with only clients, vendors, staff members, and other setup data.

The Export/Import X utility recreates a Clients & Profits X database by exporting the data into standard text files, then importing the text files into a new database. The export/import process is time-consuming, therefore, start the process at a time when you can let it run without interruptions. Keep in mind, if the process is interrupted for any reason you must restore a backup copy and begin the process again.

The Export/Import utility is only available from the Clients & Profits helpdesk. To obtain a copy email the helpdesk at helpdesk@cnp-x.com

■ If you're using Clients & Profits 4.x, do not use the Export/Import 5.x utility. There's a special version of the utility for old versions of Clients & Profits. Also, it won't work on Clients & Profits ez/4/PLUS databases.

- The process CANNOT be done while users are using Clients & Profits.
- Don't run the Export/Import utility across the network. Use a local hard disk (that is, one that is inside the computer running the utility).
- Make a backup copy of the database first.
- Let the Export/Import utility run overnight--but make sure nothing will interfere (such as backups, screen savers, Applescripts, etc.)
- If your database was damaged by a hard disk crash or other technical problems, be sure to find and solve the problems before recreating the database. If the problems that caused the corruption aren't fixed, your database is likely to crash in the future.

Preparing for an export/import

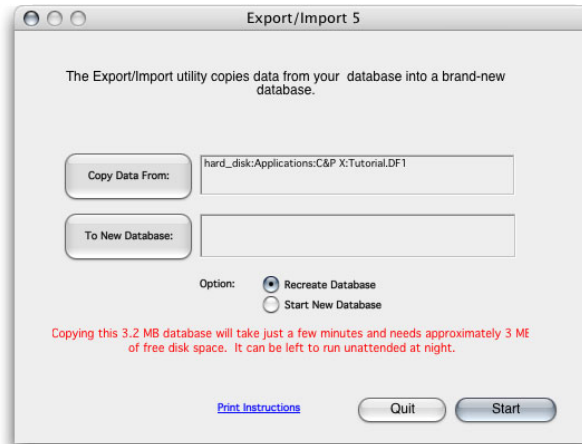
The entire Export/Import process can take anywhere from a few hours to a few days, depending on the size of your database and the speed of the computer. The faster the computer, the sooner the database will be recreated.

- 1 Choose a fast Windows 2000/XP or Power Mac G5 computer that can be dedicated to the Export/Import utility for a short time.

This computer can't be used for anything else while the Export/Import utility is working.

■ Don't export/import the database on the same computer on which the problems occurred. The same problem that caused the corruption is likely to damage the new database.

- 2 Check the computer's hard drive for free space. It should have enough space available to hold two copies of the database.



The Export/Import can be run entirely unattended. The process is completely automated so, once started, it needs no supervision until it has finished. Once completed, a report will print summarizing the records successfully imported.

The computer's hard drive needs the space to hold the damaged database and the new database. For example: a 40mb database needs 80mb of free space to run the export/import.

- 3 Use a disk first-aid tool (e.g., Norton Disk Doctor, Disk First Aid, Disk Utility, etc.) to check the computer's hard disk for disk fragmentation or other drive problems that could interfere with the export/import process.

If the computer's hard disk is too full, or doesn't have enough contiguous free space, the export/import procedure will take much longer. There's also a potential for the newly created database to be less-than-optimally organized. If there's any doubt about the condition of the computer's hard drive, reformatting the drive (or volume) is recommended.

■ If your database was damaged by a hard disk crash or other technical problems, it's critically important to find and solve the problems before recreating the database. If the problems that caused the corruption aren't fixed, your database is likely to crash in the future.

To recreate a database

- 1 Double-click on the Clients & Profits Player.
- 2 Choose File > Open Application.
- 3 Double-click on the Export/Import 5 utility.

DO NOT OPEN THE CLIENTS & PROFITS APPLICATION. The Export/Import utility is not a plug-in; it is a stand-alone utility that can't be run together with Clients & Profits Pro or Clients & Profits Classic.

The Database menu appears at the top of the screen.


- 4 Choose File > Change Data Files then select and open your database.
- 5 Choose Database > Recreate Databases from the Database menu, or click the Recreate Database button from the window.

Follow the prompts to start the export process. The Export/Import utility performs a "quick check" self-diagnostic on the database that checks for data corruption that could prevent data from being exported.

- If the Quick Check says the "data structure" of the database needs repair, the Export/Import won't work until the data file structure is repaired with Data Utilities (see right).

The Export/Import utility first counts the number of records in your database. Before the data is exported, a report is compiled showing the contents of your database. You'll need this report to compare the results before-and-after the export/import process is finished.

The Export/Import utility will methodically export data from each of the 60+ data files in your Clients & Profits database into tab-delimited text files, in alphabetical order by file name. The text file is imported in the new, empty database, then the **indexes** are rebuilt. You can track the export/import's progress by comparing the "importing..." progress message with the database count:



Your clients, jobs, costs, and accounting data is imported into a new database, called New Database (or NEWDATA.DF1 in Windows). When the data is completely imported, the utility will automatically quit.

- If your database is very large and contains two **segments**, the Export/

Import utility will automatically create and pre-size the second segment. The new database will be nearly the same size of the original database.

- 6 Rename the new database (which is named "New Database") then copy it back to the server, replacing the damaged database.

Be sure to use a different name than the original database to keep users from opening the wrong database. The export/import process is now complete. Your new database is ready for work.

To start a new database from an existing database

IF EXPORT/IMPORT DOESN'T WORK: Sometimes certain kinds of corruption can cause the export/import process to fail. In these cases, it may be necessary to repair the database first. The Export/Import utility has a built-in Data Utilities function, which you can use to repair the database before starting the export/import. To repair the database, choose Utilities > Check Data File then choose Check Data File option. Select the Quick Check option, then follow the prompts. Once the data structure has been repaired the export/import procedure should work without interruption.

This option creates a new database from your existing database, copying over only its clients, vendors, staff members, users, the Task Table, status codes, and its settings, options, and preferences. No jobs, costs, or accounting data is included.

- 1 Double-click on the Clients & Profits Player.
- 2 Choose File > Open Application.
- 3 Double-click on the Export/Import 5 utility.
- 4 Choose File > Change Data Files then select and open your database.
- 5 Choose Database > Start New Database from the Database menu, or click the Start New Database button from the window.

Follow the prompts to start the export process. Because only setup-related data is exported, creating a new database is very fast.

Appendix A: Database Troubleshooting FAQs

I tried to open the database, but couldn't. Instead, I just got an error message. Why? This error indicates either a problem with the file server or the database has been badly damaged. If you can't open the database, then the file server probably thinks the database is in use by another program—or is simply unreadable. Make sure everyone quits Clients & Profits, then try opening the database. If it still won't open, restart the server. If the database still won't open, replace it with a backup.

A user just told me that last week she got a 'bad pointer' error message when saving an A/P invoice. She continued working and since then everything has been fine. Should I be concerned? Yes. A *bad pointer*, *bad header* or *bad index* error should not be ignored. If errors like this are left unchecked and people keep using the damaged data, the corruption can be serious—and the database may become unusable. Data errors should always be resolved as soon as possible.

Our damaged database was recently repaired. It seemed to be working fine for a while, but now we're having problems again. Why? Without troubleshooting the cause of the damage, repairing the database solves only half of the problem. Whatever damaged the database will continue to affect it indefinitely—it'll be an endless cycle of data problems and repairs. See *Appendix B: Tips for Preventing Data Damage*, on page 14, for suggestions.

We're trying to repair our damaged database using Database Utilities, but it's been running for a very long time. Is this normal? Possibly not. The repair tool in Database Utilities typically takes a long time to run—about 15 minutes/mb (speed varies based on hardware). But if the repair function is stuck on a specific data file (i.e., JOBS, AR, etc.) for many hours, it's probably frozen. (Note: the COSTS file takes longest to repair since it is usually the largest.) This can happen if the database is being repaired on a fragmented or otherwise problematic hard drive. If so, check the computer's hard drive for bad blocks, fragmentation, etc.—or try a different computer—and run the repair again on a fresh copy of the database.

We're constantly verifying/recovering jobs, tasks, costs, etc. Why does it only seem to correct problems temporarily? That's because the true cause of the data problems isn't being found and fixed. Unless your system crashes while Clients & Profits is saving or posting, you should never have to verify/recover anything. But if your database is damaged, saving and posting won't work correctly every time—that's the clue that there's real trouble somewhere in the system.

How can I really know what damaged our database? Be sure to talk to the users who were first affected by a particular error—find out exactly what they were doing just before the error occurred (i.e., posting, saving, entering data, printing, etc.). Then look for any changes in the hardware or system software configuration, since something may have been installed or set up incorrectly. Use these facts to work backwards from the time the error happened, then look for patterns. You may not find the answer right away, but you'll eventually build a history that will tell the story.

Our computer consultant says Clients & Profits must have damaged the database, because he says he's never had problem with files from other programs. If your consultant may not have experience supporting multi-user databases, then they're not going to understand the nature of relational databases like Clients & Profits. Because it is multi-user, Clients & Profits is distinctly different than commonly-used programs like Excel, Quark, Word, or Photoshop. There probably aren't many other programs where the same document is used day after day for years by many users at the same time. In most productivity applications, documents are created, save, then open again infrequently—if ever. If these documents were to be damaged by a server's hard drive, you probably won't even know it. Also, the size of the Clients & Profits database is also much larger than the usual Excel or Word file. The sheer size of the database means it's a bigger target for hard disk problems and other errors. That's why you're more likely to see problems with programs like Clients & Profits.

We've checked everything and there's nothing wrong with our hardware. What now? First, look for user errors—primarily the way they handle padlocks, system freezes, and end-of-day shutdowns. A user who shuts off or restarts their computer while Clients & Profits is still running can damage the database. Restarting or shutting down the file server while users are still running Clients & Profits can damage the database, too. For a list of other common causes of data corruption, see page 15.

Are we outgrowing Clients & Profits? That's unlikely. The software is being used every day in thousands of shops. The C&P database can store data for several years, depending on the shop's volume. As long as the server is adequate to handle the work load, it's difficult to outgrow Clients & Profits. For shops with more than 50 users, upgrading to C&P SQL will provide an Oracle or Microsoft SQL Server based systems that's more scalable and fault-tolerant.

Our consultant has run Norton on our file server and says its ok. What now? Disk tools like Norton Utilities check for specific problems with hard disks, but they're not foolproof. They won't find intermittent problems caused by bad RAM, network adapters, and other upcoming hardware failures. They also won't tell you if the data is being damaged by user actions, which is more common than not. The only way to determine the true cause of data corruption is to track each incident.

How do we know that Clients & Profits isn't damaging the data? It isn't likely, because Clients & Profits doesn't tell the file server how or where to save data—which is when corruption occurs. Clients & Profits simply sends the data to the server and says "save this" and relies on the server to do its job.

Will installing the latest update fix the damaged data? No. Databases are damaged when something irregular happens when changes to records are being saved. New updates to Clients & Profits always contain bug fixes that keep bad data from occurring, but that's unrelated to database damage. Since damaged data is caused by system and/or hardware problems, any C&P update won't make a difference.

Why can't Clients & Profits prevent the damage to the database in the first place? It's a performance issue. Any kind of data verification would substantially slow C&P down whenever anything was saved, affecting users more often. Under normal circumstances, the database engine in Clients & Profits is dependable and predictable, so the extra validation isn't justified.

Why can't Clients & Profits fix any data damage automatically? That would be risky. Repairing a database is a single-user function, so Clients & Profits couldn't even start fixing an error if it found one until everyone had quit for the day. Also, a database should never be repaired before a backup is made—just in case the computer running the repair crashes.

Should we purge data to avoid adding a second segment? No. The size of the database or the number of segments has little to do with its performance. Smaller databases aren't more reliable than larger ones; they're just easier to store and back up. Purging won't shrink the database anyway (it just doesn't grow as fast).

I've download and installed all of the patches from the C&P Patches web page, but why are we still getting damaged data? Patches only fix specific bugs in C&P in between software updates. Damaged data is caused by hardware or user errors, so can't be prevented by software updates or patches.

We're using C&P Classic. Will upgrading to Clients & Profits Pro solve our data problems? No, because the Clients & Profits software itself doesn't affect the database's integrity. Clients & Profits Classic and Clients & Profits Pro share the same underlying structure and programming, so any data problems you're having now will happen no matter which versions you're using.

Our database is frequently damaged, but we only get notified by the Quick Check once in a while. Why? The Quick Check runs only when a manager-level user opens the database. If you have only one Clients & Profits system manager, this means the Quick Check will only run at most once a day. If your system manager is a part-time consultant or infrequent user, the database will be quick-checked even less often. Making more users "managers" would cause the database to self-check more frequently. Or, the stand-alone QuickChecker utility can be used to constantly monitor the database.

Why doesn't the Quick Check validate the data every time any user (not just a manager) opens the database? While the Quick Check is relatively fast, it still takes a few seconds to run. Quick-checking the database every time any user opens it would noticeably slow everyone down.

We've verified/recovered the database, but it's still damaged. Why? That's because the verify/recover utility only checks the content of the data, fixing incorrect totals and other bad data problems. It has nothing to do with the structure of the database, which is where data corruption happens. Repairing a database has no effect on the content of data itself.

What's the difference between verifying, reindexing, and repairing?

Verifying fixes bad data (e.g., incorrect account balances, job totals, etc.) but has no effect on the structure of the database. *Reindexing* fixes the index tables inside the database that Clients & Profits uses to find records. *Repairing* checks and fixes the structure that holds the database together, like the frame of a building.

What's the difference between the Quick Check and the Database Utilities?

The Quick Check checks for damaged indexes, headers, and pointers that store the data in the database. The repair tools in Database Utilities do everything the Quick Check does, plus checks the integrity of every field in every record in the database—that's why repairing takes so much longer.

Are databases less reliable as they grow larger? No. Since data corruption is caused by system, user, and/or hardware problems, a small database (i.e., less than 20mb) is as reliable as a large (over 100mb) database. However, larger databases are more vulnerable to disk fragmentation, bad drive sectors, and other hardware errors. Also, as more people use the database there's a greater chance of user error.

Will purging data from prior years prevent damaged data? No. Purging data only keeps the database from growing larger as quickly.

What does the "data files full" error message mean? It could mean (1) the server's hard disk is out of free space, or (2) the database (or segment) has reached 256mb. In any case, new data is lost because there's no place to store it. If repairing can't fix it, you'll need to restore a backup.

I'm repairing a file that has 187,000 records, but it's taking a long time -- the progress count is in the millions. Should I stop the repair process? Yes. A count that high indicates that the repair tool is stuck on a particular data file because it is seriously corrupted (Remember that the progress message isn't actually counting individual records; instead, it's an internal count.) You'll need to export/import to recover the data.

How often does the database need to be reorganized? Reorganization is needed only after installing a Clients & Profits update—and only if there's a change to the new version's files and fields.

Even if we get a Quick Check error, can we work in other areas of Clients & Profits that aren't damaged? Yes, but there's nothing to prevent someone from adding or editing records in the damaged data files, which will make them worse. It's always best to repair corrupted data as soon as possible.

How can I know if records were lost during a repair? The Database Utilities window shows the record counts for each data file. Print the data file list before and after the repair, then compare the record counts. If the "after" record count is lower, then you've lost data in those files. There's no way to know which specific records were lost.

Appendix B: 15 Tips for Preventing Damaged Data

1) *Don't shut down any computer that's running Clients & Profits, especially if the cursor becomes a padlock.* Getting a "padlock cursor" means the record you're saving is in use by another user. When the record is available, the padlock will disappear and your changes will be saved. The only safe way to quit C&P is to choose File > Quit. This logs users out of the database on the file server.

2) *Don't restart or shut down the file server while people are using Clients & Profits.* Anything users are editing or saving will be lost, and the data files they were using can be corrupted.

3) *Use Database Utilities to pre-size the database.* Pre-sizing reserves a chunk of space on the server hard drive just for the C&P database. By allocating this space, the server won't fragment the sectors around the database. The file size will appear very large, but it's simply free space for new work (see page 18).

4) *Use a dedicated file server for C&P that's optimized for databases.* Keeping the database on a separate admin file server means C&P users won't be slowed down by creatives copying graphic files. Also disable any screen savers or energy savers.

5) *Use a UPS to protect the database from power problems.* Unstable electrical power can cause data loss in unpredictable ways. A UPS replaces the surge suppressor to provide power that's free from surges, spikes, and dropouts. The UPS's battery gives you enough time to backup the data.

6) *Use a RAID system on the file server.* RAID subsystems ensure that data is saved more reliably (and often faster) in high-volume environments.

7) *Watch for disk fragmentation.* Disk fragmentation is a common cause of corruption, since it causes data to be saved erratically. Bad disk blocks also decrease the stability of the database. Tools like Norton Utilities should be used weekly to find and repair disk problems before they become serious.

8) *Don't let the server's hard drive get too full.* A lack of contiguous free space on the drive can cause the end of the database to be corrupted.

9) *Keep the server room cool.* Extreme heat can cause computer hardware (especially hard drives) to overheat, leading to data corruption.

10) *Check the database's size monthly.* Track the database's growth to keep from running out of space in the segment (it can't exceed 256mb). (See page 18-19)

11) *Keep up with new service packs, patches, and other fixes to the system software.* Usually these updates improve performance, but not every time. That's why it's important to test them first,—and be sure to have a back-up plan in case of incompatibility. You may need to quickly revert to the previous system.

12) *Look for bad wiring and other LAN problems.* Loose, frayed, or pinched cabling can cause intermittent data loss. Replace any old CAT 3 cabling with CAT 5 to prevent packet loss.

13) *Upgrade the server as the volume of work increases.* If the server is stressed, users will see more padlocks and delays, increasing the possibility that they'll reboot their computers thinking they're frozen or have crashed.

14) *Use a reliable backup system—every night.* Use a high-capacity media (e.g. Jaz, DAT, CD-R, etc.) with high-quality backup software (not Stuffit or WinZip) to make dependable copies of your C&P database. Remember to give the designated person enough time to backup the database every day—even if it means shutting down the network

early. Be sure everyone quits C&P before backing up the database. Rotate your tapes for redundancy. And test the backups periodically by restoring a copy of the database.

TIPS FOR PREVENTING DAMAGED DATA

- 1 Don't shut down any computer (or force C&P to quit) while Clients & Profits is running—especially if you get a padlock cursor.
- 2 Don't restart or shut down the server while people are using C&P
- 3 Use Database Utilities to pre-size the database
- 4 Keep the C&P database on a dedicated file server, that's optimized for databases.
- 5 Use a large UPS to protect from power problems
- 6 Use a RAID configuration on the file server
- 7 Watch for disk fragmentation or other drive problems
- 8 Don't let the server's hard drive get too full (always check for free space available)
- 9 Keep the server room cool
- 10 Check the database's size monthly
- 11 Keep up with all new service packs, patches, and other fixes to the system software
- 12 Look out for wiring and other LAN problems
- 13 Upgrade the server as work load increases
- 14 Back up the database nightly—and verify that it works

Appendix C: Common Causes of Database Damage

Cause	Description
<p>Users get padlock cursor when saving or posting, turn off computer</p> <p>Or,</p> <p>a user's computer freezes, gets restarted</p>	<p>This is one of the most common causes of data damage, because users misunderstand the meaning of a "padlock" cursor. During data-entry or saving, the "pointer" cursor becomes a "padlock" when the same file is use by some other user—whoever got it first. This is a natural feature of record-locking multi-user databases, and occurs to ensure that two people don't edit the same record at the same time. The padlock cursor disappears automatically as soon as the person who's editing the record saves their changes, usually within seconds. But there's no way to know how long the user will keep it locked while making their changes. If the user with the padlock cursor thinks their computer has crashed and restarts it, the files being updated can be corrupted. To prevent data damage, every user should know not to restart their computer while saving. Padlocks can be cancelled by pressing command-period (on a Mac) or control-break (on a PC).</p>
<p>Disk fragmentation on the server's hard drive</p>	<p>Disk fragmentation can cause the file server to save data in irregular locations on the hard drive, making it more difficult for Clients & Profits to efficient to find the data when its requested by users. Over time, sections of the C&P database end up being spread all over the hard disk, making it vulnerable to disk directory errors.</p>
<p>Power problems (surges, failures, brownouts, etc.)</p>	<p>Power failures can cause the corruption if data is being saved in the C&P database incompletely (it's as if all of the current users restarted their computers; see above). Power surges and brownouts can cause data to be saved erratically in the Clients & Profits database, making some records partially unreadable or otherwise unusable.</p>
<p>Bad RAM, network adapter, or other hardware on the server</p>	<p>Faulty hardware, defective network adapters, and other equipment problems can cause data to be saved incorrectly into the Clients & Profits database, making those records unreadable. Since hardware problems like this start sporadically, then gradually get worse, you'll see an increasing amount of data corruption until the problem is discovered and fixed. The best way to troubleshoot hardware problems is to selectively replace one component at a time until the faulty one is found. These are the toughest problems to troubleshoot.</p>
<p>Poor LAN wiring, loose connectors, electrical interference from heavy equipment</p>	<p>Low-quality network cables can cause signal loss, which forces the workstation to resend the same data several times before it is successful. Loose LAN connections can cause erratic communication with the server (which appear to the user as a padlock cursor). Frayed or pinched wiring, as well as cables that are walked on, can cause problems over time as the cables degrade. Unshielded cable laid next to elevators, large copiers, and other electrical equipment can cause intermittment communication errors, which can make a user think their computer has frozen or crashed—so they reboot it, breaking their server connection to the database.</p>
<p>Faulty system software, software incompatibilities, driver conflicts; excessive heat in the server room; disable opportunistic locking in WinNT</p>	<p>On the server side, fault system software, incompatibilities, and conflicts can cause crashes that suddenly disconnect users from the C&P database. The data files and header blocks being updated when the server crashes can then become corrupted. If the server's hard drives become too hot, data possibly won't be saved correctly. In high-volume shops, disabling "opportunistic locking" on the Windows 2000/2003 file server will better tune the server for databases (see page 21).</p>
<p>Too much server traffic; server caches set too high, so small data packets aren't written to disk</p>	<p>In shops where one server is being shared for both graphics and administration, the high traffic generated by moving graphic files can overwhelm the server and cause performance problems for Clients & Profits users. If the server slows down so much that C&P users think they've frozen, they're likely reboot their computers—breaking their connection with the C&P database, which can damage it. It's highly recommended that admin have its own file server.</p>

Appendix D: Inside the Clients & Profits Database

The Clients & Profits database contains a set of data records corresponding to one or more file formats (called **data files**), along with a set of indexes for those records. Clients & Profits stores its data files, field names, file connections (i.e., relations or joins) in a **data dictionary** inside the database.

Unlike relational databases like FoxPro and Access, Clients & Profits keeps all of its data files, indexes, etc. in a single physical file on the server's hard disk. The file has a name ending in .DF1 and looks like this:



The Clients & Profits database stores records for the 50+ data files used by the software to track clients, jobs, costs, etc. A list of the data files used by Clients & Profits can be found in the Database Utilities window (see right, or choose Setup > Utilities > Database Utilities) or from the Data Dictionary on the Clients & Profits web site: www.clientsandprofits.com/support/support/data_dictionary/data_dict.html

Each data file (e.g., JOBS, COSTS, etc.) creates a **slot** in the database. Clients & Profits also reserves some slots as **system slots**.

Database segments supports large amounts of data

A Clients & Profits database may comprise up to 10 separate physical document files, called **segments**. Each individual segment has a maximum size of 256 mb, and is identified with a special file extension (i.e., .DF2, .DF3, .DF4, etc.). Clients & Profits treats these segments as one continuous database. Clients & Profits automatically expands the first segment (i.e., .DF1) of the database as its needs more space by adding unused **disk blocks** to it. This happens behind-the-scenes as users add new records to the database. When the database nears 256 megabytes, an additional segment (which will have the file extension .DF2) must be created. These additional segments can't be created or expanded automatically, so the Database Utilities window has a special "Change Database Size" function that handles them.

■ **Always check the database size.** It's important to regularly check the size of your Clients & Profits database, especially once it exceeds 200 megabytes. It's very easy to neglect the size until the 256 mb limit is reached, which triggers the "data files full" error -- and can cause damage to the database. It's easy to see the database size: it appears at the bottom of the Introduce Yourself window every time someone starts Clients &

Profits. The Database Utilities window also has a function that can graph the database's growth over time, helping you plan when to add the next segment.

Slots, blocks, records, indexes, and pointers

The Clients & Profits database is composed of slots that contain the data from the program's various files, like jobs, clients, costs, and G/L entries. Each slot contains **blocks** that store bits of data from these files.

A single block stores 512 characters -- that's all. So a large database can have tens of thousands of blocks, each storing 512 characters (i.e., bytes).

File Name	Description	Records	Size	Status
ACCESS_PRIVS	Staff access privs to clients	27	kb	--
ACCOUNTS	G/L accounts (i.e., Chart of Accounts)	191	kb	--
ACCT_PERIODS	Prefs - accounting periods	1	kb	--
ADDRESSES	Client mail/ship addresses		kb	--
AP	A/P (header)	2,260	kb	--
AR	A/R (header)	1,158	kb	--
ARCHIVED_JOBS	Old job storage		kb	--
AR_ITEMS	A/R billing amounts (i.e., line items)	7,058	kb	--
ASSETS_FILES	Asset Manager - elements		kb	--
ASSETS_PREFS	Asset Manager - prefs	1	kb	--
ASSETS_TYPES	Asset Manager - file types		kb	--
AVAIL	Staff availability (for job scheduling)		kb	--
BILLING_INFO	Jobs - corporate accounting info	10	kb	--
CALENDAR	Company-wide calendar (i.e., cash flow, media)		kb	--
CHANGE_ITEMS	Job change orders - line items	7	kb	--
CHANGE_ORDER	Job change orders - header	218	kb	--
CHECKS	Checkbook (header)	1,298	kb	--
CHECK_ITEMS	Checkbook - line items	2,389	kb	--
CLIENTS	Client accounts	57	kb	--
CONTACTS	Client/vendor contacts	57	kb	--
COSTS	Job costs (i.e., AP line items, time, expenses, etc.)	14,420	kb	--
CREDIT_CARD_PREFS	Credit cards prefs - accounts + merchants		kb	--
DEPTS	Departments	7	kb	--

The Database Utilities window in Clients & Profits offers system managers a single console to monitor the database's growth, check its stability, and change its size. It also includes tools for exporting and importing data.

There are five types of blocks in the C&P database:

- Data blocks
- Index blocks
- File header blocks
- Free blocks
- The "master" block

The block that stores the information for a database record is a **data block**. Depending on the size of an individual record, one data block may be large enough to store several records from a data file (such as time entries, which are small), or Clients & Profits might need several data blocks to store just one record (such as purchase orders, which are large).

Index blocks store index values. Every data file has special index fields which are used for finding records, such as the client number, job number, task, etc. Clients & Profits indexes the data values in a tree structure. Every record added, changed, or deleted rearranges the tree.

Free blocks are unused blocks in the data file. Clients & Profits keeps track of the free blocks in a data file in a *free chain*. Clients & Profits updates the

free chain whenever it deletes a record from a specific data file by marking the block or partial block (depending on the size of the record) that the deleted record occupied as being available for new records.

Clients & Profits maintains a **file header block** for each data file in the database. File header blocks store information about the attributes of each field in a data file. This lets C&P determine whether a particular data file needs to be reorganized with each new update. The file header block also maintains a pointer to the first block in the data chain for that file, as well as a pointer to the first and last block of all indexes for that file (and other information).

Block 0 is the **master block** for the database, which contains critical information about the database. It holds the total block count for its own segment of the database and, if there is more than one segment in the database, also holds information about the sizes of the other segments. Block 0 also stores a **pointer** in the first block in the **data set free chain**, a directory of all files in the database, and a point to each of the file header blocks.

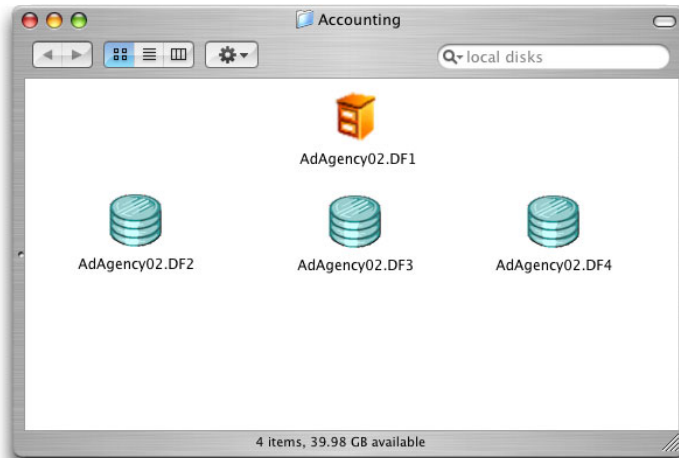
The pointer is like a street address. It tells Clients & Profits where a specific record (for example, a job ticket) is stored in the database. When a user finds a job ticket, Clients & Profits requests the record from file server using its address (i.e., by pointing to the block in the database). Clients & Profits doesn't actually know where the data is stored physically on the file server's hard disk. All it knows is the slot, block, and pointer—the unique address for each record in the database.

Managing multiple databases

A separate **database** file is created for each database you'll use, such as the tutorial database, the sample database, and your shop's database. Each database file works like any other application file; it is opened, moved, and backed up like any word processor or spreadsheet document. If your accounting department manages the books for several different companies, you'll probably have several different database files on your server's hard disk. Each database will have a unique file name, but will always end in .DF1.

How Clients & Profits uses the Quick Check to look for trouble

The "quick check" function looks for early data corruption in the Clients & Profits database. It runs automatically each time a manager-level user clicks OK in the Introduce Yourself window.



The Clients & Profits database can contain 10 segments, containing up to 256mb of clients, jobs, costs, and other data each. Each database segment has a special file name extension that identifies its order in the database structure.

The Quick Check examines the status of the database by reading only the internal headers, indexes, and pointers (not the records themselves, to save time). These records indicate corruption caused by either hardware, network, or other system failures. It doesn't attempt to systematically check the entire database for problems; that's the role of the Check Data Files function in Database Utilities. For this reason, running a quick check only takes a few seconds.

Because it is optimized for speed, the Quick Check doesn't find every kind of database damaged. However, it is very effective at finding data corruption early,

before it has a chance to affect the accuracy of the data. If the Quick Check function discovers some corruption, it displays what it found in a log window. If no one else is using the C&P database, clicking the Repair button will attempt to fix the corruption immediately. **BE SURE TO BACK UP THE DATABASE BEFORE IT IS REPAIRED** since any kind of system crash during a repair function will make the database unusable.

Since the Quick Check runs every time a manager opens the database, it is likely to find data corruption within minutes or hours of its occurrence. This time frame is important, because system managers can use it to track down the cause of the corruption while the day's events are still fresh in users' minds. In shops with many manager-level users, this means the quick check is run dozens of times each day.

Appendix E: Database Size

The Clients & Profits database has a maximum size of 2.56 gigabytes, which is comprised of 10 segments of 256 megabytes each. The first database segment, which was created when you first started Clients & Profits, increases in size automatically as new records are added. As the staff learns more about Clients & Profits, the volume of work added increases. You'll see the database's growth rate quicken after the first six months. How fast your database grows is also a factor of your shop's growth—picking up several new clients will add an unexpected growth spurt to the database. It's easy to lose track of the database's size, which can suddenly run out of room—causing data to be lost. So tracking the database's size is a critical role for the C&P system manager.

To check the database size

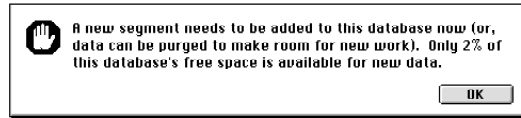
There are two ways to check the size of the database: (1) from the Introduce Yourself window and (2) from Database Utilities.

The Introduce Yourself window, which you'll see every time you open the C&P database, shows the database size at the bottom of the window. By noticing the database's size from week to week you'll get a good idea of its growth rate.

The database's current size is also shown in the Database Utilities window (see above right). The Database Utilities shows two sizes: file size and data size. The file size is the size of the database file on the file server (i.e., its physical or "logical" size). In contrast, the data size shows the actual size of the records contained inside the database file. If you don't pre-size or purge your database, these two sizes will be nearly identical. Pre-sizing the database changes the file size, since it is the actual physical size of the file on the server's hard disk. As data is added to the database, its data size increases. The data size is more important because it determines whether the database is too full—when the data size reaches 250 megabytes, it's time to add another segment or else it'll run out of space.

■ For an explanation of database segments, see page 16-17.

As your database reaches 250mb, you'll be prompted to add a segment.



This is a very important step—if a new segment isn't created, the database will run out of free space (you'll see a "data files full" error message) and data will be lost (if C&P mail is enabled, a reminder will be e-mailed to the system manager each time a manager-level user starts Clients & Profits, so the issue should be hard to ignore.)

■ The database's growth rate can be graphed in Database Utilities by clicking the button. The database's size is recorded each time this button is clicked. By regularly viewing the database graph, you'll be able to easily track the growth rate and better predict its future size.

To pre-size the database

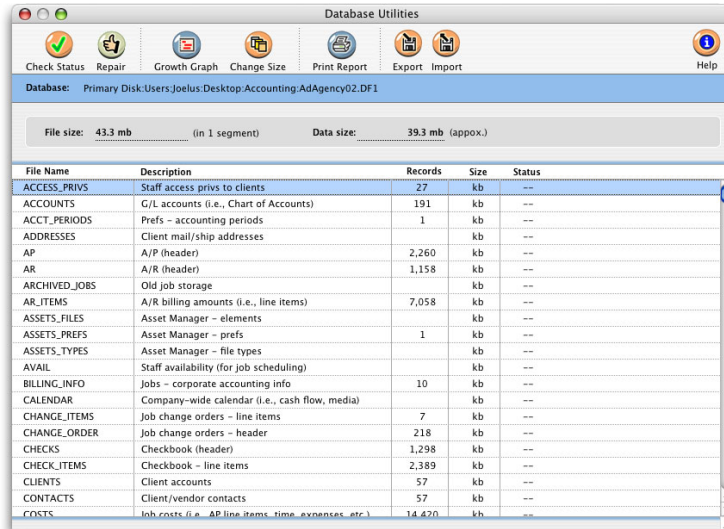
It isn't necessary to pre-size your database with Clients & Profits version 3.0 or later. Your database will expand automatically as new data is entered, but only for the first segment. But in high-volume environments (for example, more than ten simultaneous users) it is still more efficient to set the database size in advance.

By pre-sizing you're allocating a fixed amount of disk space for the database. The database's file size on the server will be larger than the amount of data that it actually contains. When new work is added, the database size won't grow—instead, the new data simply fills up the unused space. In this way pre-sizing the database helps prevent disk-based fragmentation.

Pre-sizing is a single-user function. It won't work if other users have the database open. Also, it is best to pre-size the database on a local copy, then when the process has completed, move it back to the server.

1 Choose Setup > Utilities > Database Utilities.

2 Click the (i.e., Change Database Size) button or choose File > Change Database Size.



The Change Database Size dialog box opens, prompting you to enter the database's new size. The database size can be increased in 10% increments by click the **+10%** button or by entering the database's new size (in megabytes) in the size field, up to 256 megabytes. If your database is growing relatively slowly (e.g., less than 50 mb per year) there's no advantage to pre-sizing it to 256mb—it'll just be a more difficult to copy to zips, etc.. If your database is growing faster, or you have more than ten simultaneous users, pre-sizing the database to its maximum size of 256mb is recommended to prevent fragmentation.

3 Enter the database's new size, then click OK.

Extending the size of the database will take only a few minutes, depending on how large it is being expanded. Once completed, the database is ready for new work.

The database size can be extended by either entering its new size, in megabytes, or by clicking the +10% button one or more times. An additional segment can only be added once the database has reached 250 megabytes.

To add an additional segment

A new segment can't be added unless the database is already 250 megabytes or larger. This limitation prevents a second segment from being added accidentally (which is important since it can't be undone). If someone tries to add a new segment too soon, they'll see this error:




Like pre-sizing, adding a new segment is a single-user function. It won't work if other users are using the database. Also, don't add a new database segment if there is a lot of network traffic (because if the file server crashes, the database and its new segment will be unusable and a backup copy will have to be restored).

■ First, be sure the database is properly named and has the *.DF1* file extension. The database name shouldn't have quotes, dashes, slashes,

periods, spaces, or other special characters (underscores are ok). A good name should include the company name, something that describes it as the accounting database, and the fiscal year (e.g., *CP_ACCT99.DF1*, *AGENCY_ACCT99*, etc.).

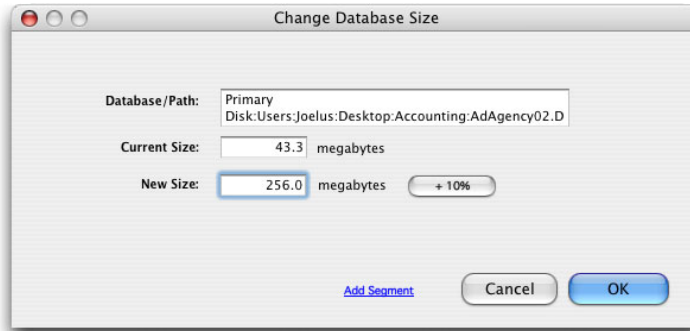
1 Choose Setup > Utilities > Database Utilities.

2 Click the  (i.e., Change Database Size) button or choose File > Change Database Size.

3 Click the Add Segment button.

4 Click OK.

The new segment is automatically named using the database's name, but ending with the *.DF2* extension. New segments are pre-sized to 100 megabytes.



■ Unlike the first database segment, additional segments aren't automatically extended as new data is added. This means it's important that the system manager track the database's free space (i.e., the data size) from month to month to prevent the database from becoming full. As the data size nears 100mb, you'll need to use the Change Database Size utility to extend the database size.

■ DO NOT MOVE, DELETE OR RENAME A DATABASE SEGMENT. Once an addition segment is created, moving, deleting, or renaming it will cause the entire database to be unusable.

■ Always back up and restore a database and its segments together as a group. If different versions of a database and its segments are mixed up, the entire database will be unusable.

■ Adding three or more segments is possible, but not recommended. For databases over 500 megabytes, upgrading to Clients & Profits SQL is recommended. For more information, contact the C&P Helpdesk at (800) 521-2166.

Appendix F: Troubleshooting Tips

While the Clients & Profits relational database is generally stable and reliable, it does have an essential weak point in high-volume environments: there is no server-side error-checking of the data, unlike Oracle and SQL Server-based databases (such as the one used in Clients & Profits SQL). This means Clients & Profits relies on the user's PC to write data to the file server. Any problems on the client side, the network, or with the file server could cause lost or incomplete data, possibly corrupting the database.

The impression that the network is OK is not always correct, especially in high-traffic situations. In these cases, many times collisions, lost packets, and intermittent connectivity become more evident. More often than not, the cause of the corrupted datafile is some network issue. There are several steps a system manager can take to better "bulletproof" their network and user systems. Consider the following tips if you're regularly experiencing problems like these:

- Frequently occurring "update or delete command without a read/write file" error messages
- Frequent padlocks when starting or quitting Clients & Profits
- Changes are frequently not saved when editing jobs, estimates, traffic, and schedules
- The database becomes damaged often (e.g., weekly), especially when many users are working together
- The users' PCs frequently freeze or crash when saving

Use caution with "optimization" utilities: Several popular system optimization programs (e.g., First Aid, etc.) suggest to the user that disabling "write-back caching" is bad, and will often re-enable it. Since few computer technicians (and even fewer users) are familiar with multi-user relational databases like Clients & Profits, be sure that they're aware that write-back caching can be harmful.

To disable SMB in OS X

Apple's implementation of SMB (the Windows network protocol) is flawed in that it does not properly support byte range locking. Connecting via SMB will cause problems with your database, so we recommend turning it off completely on your OS X machines.

To do this, follow these steps:

- 1 Open the directory access utility (by default Macintosh HD > Applications > Utilities > Directory Access).
- 2 Click the lock button in the lower left hand corner to make changes.
- 3 Uncheck SMB under "Services" and click Apply.
- 4 Click the lock button in the lower left hand corner again to lock the changes.


- 5 Close all of your programs and restart your Mac.

Be sure to repeat these steps for every Mac that uses Clients & Profits.

To disable the "write-behind caching" in Windows 2000/XP

Write-behind caching is a technique employed by Windows to improve a PC's performance. Write-behind caching attempts to optimize how the hard drive saves data from the computer's memory. This type of caching stores information that needs to be written to the hard disk and sends it when the system is idle or after a certain amount of time has elapsed. This is a built-in feature of Windows 2000/XP and makes the PC slightly faster by essentially not double-checking that the data from memory was saved entirely to the drive.

Single-user, document-based applications, such as Microsoft Word, Excel, and PowerPoint, are saved in such a way that the write-behind cache isn't needed. Database applications (especially multi-user programs like Clients & Profits) expect that the data in memory is saved completely to the hard disk. When the write-behind caching is enabled, it's possible for user's data to be saved incompletely (or not at all). By disabling the write-behind cache, the PC saves data slightly slower--but much more reliably. In cases where unexplained data problems are occurring regularly, the write-behind caching should be disabled.

- 1 Click on the **Start** button and choose Settings > Control Panel > System.
- 2 Select the **Hardware** tab.
- 3 Click on the **Device Manager** button.
- 4 Click on the  symbol next to **Disk drives**.
- 5 Right-click on your PC's main drive and choose **Properties**.
- 6 Deselect the option **Write cache enabled**.

- For Windows XP, click on the **Policies** tab instead then deselect the **Enable write caching on the disk** option.

Be sure to repeat these steps for every Windows-based PC that uses Clients & Profits, including your file server.

To disable "opportunistic locking" on the Windows 2000 server

Opportunistic locking is a technique used by Windows 2000 to increase the file server's performance by performing read-ahead, write-behind, and record lock caching. Opportunistic locking is a feature on Windows NT 3.51 and later versions. These instructions work with Windows 2000.






It works like this: if one client is accessing a block range in a file, that range is marked for opportunistic locking and the client can perform read-ahead, write-behind, and lock caching. If another user attempts to write to that block range, the opportunistic locking has to be switched off for the previous client and the data needs to be synchronized with the file server before the second user can access the range.

The architecture of multi-user relational databases conflicts with opportunistic locking, since databases involve nearly constant reading, writing, and locking of the data file. Disabling opportunistic locking will improve the performance and reliability of the C&P database, but makes the server slightly less responsive for non-Clients & Profits users.

■ For more information on opportunistic locking in Windows 2000, please refer to these URLs:

<http://support.microsoft.com/default.aspx?scid=KB;EN-US;Q296264>

WARNING: Opportunistic locking is disabled by editing the server's registry information. Incorrectly editing the Windows 2000 registry can cause server problems, including a loss of service. It should be performed only by an experienced Windows 2000 technician.

- 1 Choose Start > Run, enter the word regedit then click OK.
- 2 Click on the  symbol next to **HKEY_LOCAL_MACHINE** to expand the local computer's settings.
- 3 Click on the  symbol next to **SYSTEM** to expand the system settings.
- 4 Click on the  symbol next to **CurrentControlSet**.
- 5 Click on the  symbol next to **Services**.
- 6 Click on the  symbol next to **MRXSmb**.
- 7 Click on the heading **Parameters**.
- 8 Choose Edit > New > DWORD value.
- 9 Enter the **Name** field as `OplocksDisabled`.
- 10 Press Return.
- 11 Right-click on `OplocksDisabled`, choose Modify, change the value to 1, then click OK.
- 12 Close regedit then restart the Windows 2000 server.

The changes you made take affect immediately once the server has restarted.

To optimize the Windows 2000 server for "background services"

The Windows 2000 server is optimized by default to maximize applications, not sharing files. This setting makes copying documents to and from the server faster, but isn't an efficient setting for relational databases like Clients & Profits. A simple change to the server's network configuration will improve the performance of Clients & Profits, especially on servers with lots of users.

■ The server must be restarted after changing this setting. Before making this change, be sure that no one is currently connected to the server.

- 1 Choose Start > Settings > Control Panel > System.
- 2 Click on the **Advanced** tab.
- 3 Click on the **Performance Options** button.
- 4 Set the application response setting to optimize performance for **background services**.
- 5 Click OK.
- 6 Close the System Properties and Control Panel windows.
- 7 Restart Windows 2000.

CLIENTS & PROFITS X online helpdesk web site map



www.cnp-x.com/support

