



Mining your Distribution Data

By Jason Bader

Managing Partner - The Distribution Team

What percentage of your distribution software are you currently utilizing? I ask this question to most of the groups I speak in front of. I get some really interesting answers. I always know when somebody's boss is in the room. "We use around 70 to 75 percent." While I don't want to challenge this person's concept of reality, most of us don't use 75% of the capability of the common calculator.

In fact, most distributors I work with use somewhere between 10 to 30 percent of built in functionality. After I roll out this little factoid, I pose another question, "What percent of that distribution software did you pay for?" With the investment associated with most distribution software these days, system utilization should be something that we all strive to improve.

In previous articles I have suggested that understanding the reporting functions of your distribution software is one of the quickest ways to bump up your system utilization. The ability to extract data and put it into a meaningful format can take you from average user to power user. There are several canned reports, well over 100, built into the system. While these reports are meaningful and can be interpreted to fit our needs, they rarely grab our attention and motivate us to act. Now if these standard reports came with a billy club and an air horn, we might just jump to a conclusion.

Most good distribution software gives the user the ability to create custom reporting. Unfortunately, most managers are not willing to invest the time or the energy to learn the reporting functions. Who can blame them? They have fires to put out. Actually, until very recently, most distribution reporting functions were less than intuitive. And I am being really nice here. In fact, they were a real pain in backside. This is why reporting functions were delegated to the company propeller head. Hey, I am not slinging mud here. I was the "computer guy" for more years than I care to recall. The fact of the matter is that reports were hard to design and became very frustrating for most. We now have something better.

Business Intelligence software, or BI, has been around for several years. You may be familiar with several products out there: KMC, MITS, and TLG just to name a

few. When I was first exposed to the concept in the 90s, it was commonly referred to as data mining. Data mining is the process of extracting data from your daily transaction software and storing it in an alternate database. The basic idea here is to create a warehouse of information about your business. You can store information about all aspects of your business. As long as the main transactional system captures the data, you can store it in the warehouse. Once the data is stored, break out the pick and start digging.

The real advantage here is that this data warehouse is typically stored outside of the main transactional server. While this is not a requirement, it is highly recommended. The general set up requires that you invest in another server. Great, here we go spending more money. While it may seem like this is counter intuitive to our ultimate goal of reducing cost, this one is a really good investment. The best argument for housing our data in an alternate server is processor productivity. Have you ever wondered why your distribution software package seems to come to a grinding halt in the middle of the day? This is especially true if you work in a branch location running through a fractional T-1 line or DSL. Sorry, too much geek speak. It has been my experience that these slowdowns usually occur around some reporting activity. Although software developers have given us access to reporting tools, they may be dangerous in the wrong hands. Let's just call it the super query. Someone, who will remain nameless, wants to find out how many pan head machine screws we sell to a particular salesman number below a certain margin during the last 5 days of the month. While we are at it, let's get the order number and see what else was bought on those transactions. Did I just throw out a ridiculous scenario? Perhaps, but some of you are nodding your head. This isn't out of the realm of reporting requests. Because the software systems have given us the ability to create this report, we generate this Frankenstein of data design and then push run. Can the system do it? Sure, but don't expect anything else to happen while your main transactional server grinds through the database. For those of you in the character based world, this is why you see screens refresh one agonizing letter at a time.

Get the cumbersome reporting functions out of the main transactional server. By housing the data in an alternate high speed server, you can query to your heart's content without affecting the daily transaction functions. I am not talking about spending a ton of money in hardware. Pick up a windows based server with a bunch of memory. Make sure to get plenty of processor power and RAM. This will meet your needs for several years.

One of the first steps is to figure out what data you want to pull over from the main transactional server. Of course you want it all; but that is a good way to fill your data warehouse in short order. The beauty of using one of these canned BI products is that they work with a standard set of recommendations. They suggest that you pull certain data that most users have found useful over the years. Building cubes of data is another way to describe this data extraction. If the BI product is building a sales cube, it will pull data related to sales transactions. An inventory cube will pull data from product movement related fields. If you feel that the standard cube is not covering your needs, they can be modified. Just remember, modification usually comes with a hefty price tag.

The next big decision refers to timing. How often should we extract data from the main transactional server? Most companies feel that nightly extraction is more than adequate. In many situations, weekly would do the trick. I have some clients that want to extract data every 5 minutes. They feel that sales data isn't worth much if it is 24 hours old. The problem with frequent extraction is that you really fill up your warehouse too quickly. If you want to have quick sales data, set up simple reports in the main transactional server. Leave the longer range trending to the data warehouse.

Many of my more tech savvy readers will suggest that you don't have to spend the money on these packages. Everything they do can be created in house. I agree with them. A good programmer with a background in database architecture can pull this off. But seriously, why reinvent the wheel? Why not just take the basic package and improve on it? That seems to be a more prudent use of your time and energy.

While canned programs have their limitations, their real strong suit is in the ease of use. With minimal training, managers can navigate the reporting function and even create their own customized views. Each morning, a manager could create their own little view of the business. The graphs, charts and visual images would pertain to specific areas of responsibility. Imagine a world where you don't have to sift through spreadsheets and columns to get the real answer. Am I making money or not? Believe me, this is not always the easiest question to answer. A good BI system can help.

If you looking at new distribution software, or are considering an upgrade, be sure to take a look at the business intelligence offerings. The ability to extract data, and put it into a meaningful format, will improve your system utilization dramatically. There's gold in them thar gigabytes. You just need to have the right pick.

Good Luck.