

Configuration: Discover the brain behind the brawn

By Mick Legare

It is clear that the Configuration Advisor's speed and efficiency should become an integral part of database systems design and implementation. Experiments show that in seconds, the Configuration Advisor can tune a system to perform within 10% of an expert's tuned settings and showed improved performance on some customer's environments. In one customer experiment, the Configuration Advisor was able to achieve the best-ever system performance for the customer's database, with the new configuration performing at 224.72% of the original administrator-tuned throughput. This section describes the Configuration Advisor in more detail and shows through experiments, how effective it is at configuring systems quickly.

Feature details

The objective of the Configuration Advisor is to define the database configuration parameters and memory assignments to optimize system performance. The complexity of the memory topologies for enterprise systems (including non-IBM solutions) makes these assignments a difficult task, even for skilled administrators deeply aware of their workload characteristics. The Configuration Advisor is designed on the principle that configuration choices can be made by modeling each database configuration setting as a mathematical expression that combines three sets of information. The three sets of information are:

- Your specification of the database environment (designed as a small set of basic input, generally requiring minimal skill)
- Automatically sensed system characteristics (such as number of CPUs, disks, amount of RAM, number and size of relational tables, etc.)
- Expert heuristics for database configuration, as reported by experienced database administrators and performance tuning experts

The first two of these information sets are essentially parametric, and can typically be represented as scalar values you input either through the CLP or the Configuration Advisor GUI. The Configuration Advisor is designed so that the responses collected from you make up the minimal set of characteristics that are required to effectively configure a database. The answers require neither lengthy nor detailed analysis of the database.

The interactive portion of the Configuration Advisor gathers the following information:

- Percentage of server memory to use for the database

- Workload type (online transaction processing, decision support, or mixed)
- Average number of statements per unit of work
- Estimated number of transactions per minute
- Prioritization for transaction performance and database recovery, or whether there is no priority at all
- Whether the database is currently populated with data
- Average number of connected local applications
- Average number of connected remote applications
- Isolation level (row-locking) requirements (repeatable read, read stability, cursor stability, or uncommitted read)

After stepping through these simple questions, the Configuration Advisor combines your responses with the information it automatically detects about the database server. This information is described in greater detail in a later section.

The second set of input information is obtained through the automatic detection of system resources, which includes the number of physical disks, physical memory size (RAM), number of online CPUs, number of configured CPUs, OS features, OS type and release, database size, number of tables, number of indexes, number of table spaces, name and size for each buffer pool, and the quantity of buffer pools.

The first two classes of parametric information are then combined with a set of expert heuristics to define the configuration model. The expert heuristics were collected through an ad hoc process of surveying over a dozen of DB2 UDB's leading performance experts and architects, as well as through review of published information on DB2 UDB performance tuning.

The configuration model expresses the configuration settings as a mathematical expression.

The model for the configuration settings is further divided into three distinct classes:

1. Independently modeled configuration settings, which can be modeled independent of other configuration settings
2. Dependant configuration setting, where the value of one setting affects the model of another (or perhaps co-dependency), and
3. Zero sum game relationships (such as memory for sort, caching, locking etc.) in which a fixed resource must be divided among a set of configuration parameters.

The configuration model yields a set of configuration recommendations. These recommendations are then presented to you, where you may choose to apply them. The modeling effort performed by the Configuration Advisor for these recommendations typically

requires less than ten seconds. The following configuration parameters specific to DB2 UDB are recommended and can be automatically set by the Configuration Advisor:

Database Manger Configuration

AGENT_STACK_SZ	ASLHEAPSZ	FCM_NUM_BUFFERS
FCM_NUM_RQB	INTRA_PARALLEL	MAX_QUERYDEGREE
MAXAGENTS	NUM_POOLAGENTS	NUM_INITAGENTS
PRIV_MEM_THRESH	RQRIOLBK	SHEAPTHRES

Database Configuration

APP_CTL_HEAP_SZ	APPGROUP_MEM_SZ	APPLHEAPSZ
CATALOGCACHE_SZ	CHNGPGS_THRESH	DBHEAP
DFT_DEGREE	DFT_EXTENT_SZ	DFT_PREFETCH_SZ
DFT_QUERYOPT	LOCKLIST	LOGBUFSZ
LOGFILSZ	LOGPRIMARY	LOGSECOND
MAXAPPLSMAXLOCKS	MINCOMMIT	
NUM_IOCLEANERS	NUM_IOSERVERS	
PCKCACHESZ	SOFTMAXSORTHEAP	STMTHEAP
STAT_HEAP_SZ	UTIL_HEAP_SZ	

BUFFERPOOLS

The Configuration Advisor was built for a wide range of machine sizes and was tested on machines with memory configuration ranging between 512 MB to over 500 GB. It was shown that the formulas used by the Configuration Advisor scaled very well between the very small and the very big. The graph in Figure 1 shows the Configuration Advisor's recommendations for the database configuration parameter MAXAPPLS. The Configuration Advisor ran with two sets of hints:

- 1) num_remote_apps 500 num_local_apps 0 mem_percent 80
- 2) num_remote_apps 5000 num_local_apps 5000 mem_percent 80

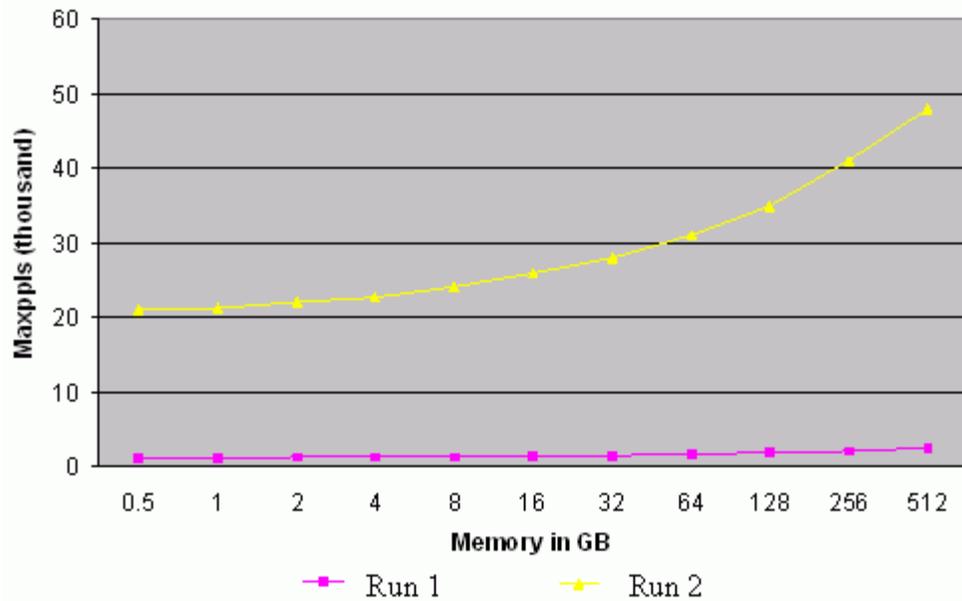


Figure 1. Maxappls recommendation (in thousand) in function of machine size (amount of memory in GB) for Configuration Advisor runs with 2 sets of hints.

The experimental results of this modeling are described in the following section.

The Configuration Advisor in action

To examine the effectiveness of the Configuration Advisor, four experiments were performed on systems running distinct workloads and environments, comparing the performance of the database after tuning by the Configuration Advisor to the system performance achieved through tuning by an expert database administrator. These experiments include a degree of inaccuracy, since the tuning performed by human administrators has variable quality, and it is impossible to assess how close this tuning is to optimal. However, what these experiments do illustrate is the degree to which the Configuration Advisor is able to tune the database system relative to the human expert. While the human administrator typically spends a number of days tuning these configuration parameters, the Configuration Advisor provides recommendations for the same set of parameters within seconds.

The four experiments included:

- **Online Transaction Processing (OLTP) industry standard workload**

This workload was run once on a 32-bit implementation, and again on a 64-bit implementation. This benchmark simulates a population of terminal operators executing transactions against a relational database. The benchmark models the transaction

environment of an order-entry environment. Two operational workloads tested on-site at two of the world's leading global investment banks. While all tests in this set of experiments were run on AIX operating systems, this was more due to general availability of test systems rather than any platform-specific constraints on the Configuration Advisor.

- **64-bit OLTP**

The 64-bit OLTP experiment was run on an RS/6000^(R) 44P Model 270 4-way 375 MHz Power3-II server, configured for 2-way processing. The system had 8 GB memory, and was running AIX 4.3.3 with DB2 UDB EE V7.2 (64-bit). The storage system on this server included 3 ServeRAID 4H (SCSI) adapters as follows: fifty-six 9 GB SCSI Disks (data), fifty-six 9 GB SCSI Disks (data), fourteen 9 GB SCSI Disks (log, backup, temp table spaces). The hand-tuned system was configured by IBM performance specialists. In this experiment the database performance after tuning with the Configuration Advisor was observed to be 93.58% of the hand-tuned system.

- **32-bit OLTP**

The 32-bit OLTP experiment was run on an RS/6000 44P Model 270 4-way 375 MHz Power3-II server, with 8 GB of memory, running AIX 4.3.3 with DB2 UDB EE V7.2 (32-bit). The storage system on this server included 3 ServeRAID 4H (SCSI) adapters as follows: twenty-eight 9 GB SCSI Disks (data), forty-two 9 GB SCSI Disks (data), fourteen 9 GB SCSI Disks (log, backup, temp table spaces). As in the 64-bit experiment, the hand-tuned system was configured by DB2 UDB performance specialists. In this experiment the database performance after tuning with the Configuration Advisor was observed to be 91.52% of the hand-tuned system

- **Global investment bank A**

Two sets of tests were conducted on production workloads. The first of these tests was conducted on a high volume OLTP test system that stress tested the database using an authentication application. The application simulated authentication operations that would result from customer accesses of the bank's products. The test environment consisted of a 6-way RS/6000 Model F80 server with each processor running at 500 MHz. The server had 4 GB of memory and was running AIX 4.3.3. Its data was spread over 14 of its 16 drives (two 9 GB drives and twelve 18 GB drives). In this experiment, the hand-tuned system was configured by the bank's database administrator. The

results of the experiment show that after the Configuration Advisor was run, the workload performed 5.62% better than the hand-tuned system.

- **Global investment bank B**

The second of the two production system tests was conducted on a system that allows customers to view their Web pages and change their preferences. These preferences are sequenced into records in a DB2 UDB database. The transaction types are mixed between selects and inserts or deletes, with 60% of the workload being selects. To provide for load balancing and for hot backup solutions, the system exploits two servers with peer-to-peer replication. The two servers had identical hardware specifications and topologies. Both were IBM RS/6000 Model 270 4-way 375 MHz servers with 4 GB of memory running AIX 4.3.3. The storage system included sixteen 18 GB drives. At the time of the experiment, the research prototype of the Configuration Advisor did not include a model for systems that deploy peer-to-peer bulk replication. As a result, the Configuration Advisor under-configured memory for locking and logging, resulting in less than optimal system performance. The administrator at the customer site decided to leave the settings for locking and logging at their pre-Configuration Advisor values and adopted all of the remaining Configuration Advisor recommendations. The result was a best-ever system performance for the database, with the new configuration performing at 224.72% of the original administrator-tuned throughput. It is interesting to note that the alteration in buffer pool size (set by the Configuration Advisor) was an important change in the configuration. While the Configuration Advisor was clearly inadequate without the planned extensions to support bulk replication, the recommendations of the Configuration Advisor were effective in combination with two adjustments by the database administrator (for locking and logging) in achieving a massive improvement in system performance.

Witness the difference

In all of these experiments, as shown in Figure 2, the Configuration Advisor was able to evaluate and compute a revised or proposed configuration in two or three seconds. The performance of the Configuration Advisor is significant when compared to the amount of time typically spent configuring large database management systems, usually on the order of one to two weeks.

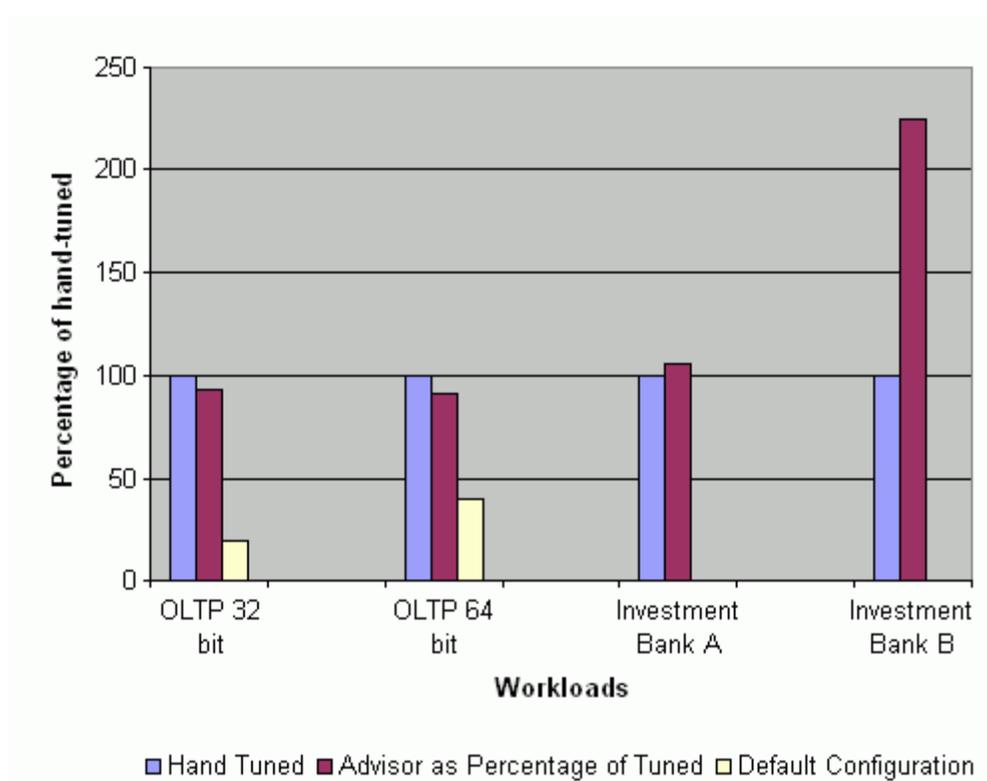


Figure 2. Database Performance Comparisons

In the two cases in which the Configuration Advisor was tested against systems configured by IBM performance experts, the resulting performance was within 10% or less of the hand-tuned result. As well, in these two cases when the Configuration Advisor selections were tested against the default configuration, the Configuration Advisor significantly outperformed the default settings. In the two cases where the Configuration Advisor was tested against databases in the field, the Configuration Advisor exceeded the hand-tuned result (though in the case of "Global investment bank B" two parameters were left at the hand-tuned setting to account for the special needs of peer-to-peer replication).

It should be noted that in a partitioned database environment, the Configuration Advisor only applies changes to the current partition. You can use the **db2_all** command to run the Configuration Advisor and apply the changes on all database partitions. Since the Configuration Advisor needs a connection to the database you would run the command as follows:

```
$ db2_all ";db2 connect to sample;db2 autoconfigure using mem_percent 80 apply db only;db2 disconnect sample"
```

With regards to buffer pools, the current recommendations are quite simple. The Configuration Advisor estimates the total amount of buffer pool needed and distributes it among all the available buffer pools. The distribution is weighted based on the initial size of each buffer pool. Since the Configuration Advisor has no knowledge about what each buffer pool contains, it considers them all equal value. If you have more than one buffer pool, we recommend that you note the Configuration Advisor's buffer pool recommendations and redistribute as needed. You should keep the total amount of buffer pool size close to the Configuration Advisor's total buffer pool recommendations since this value was derived based on other configuration elements.

The Configuration Advisor has the ability to reduce and possibly eliminate the tedious and time-consuming task of configuring a system for desired performance as is shown with the benchmarking results. These results indicate that the Configuration Advisor shows promise for providing quality recommendations on database configuration, and, in some cases, exceeding the performance quality of human-configured systems, as observed in the case of the two investment banking systems. If performance on the system is not an absolute priority, the Configuration Advisor eliminates the need for frequent manual tuning of performance related configuration parameters. If the system is required to be configured for near optimal performance, the Configuration Advisor supplies a configuration that can serve as a springboard for further fine-tuning based on specific workload characteristics. The simplicity and speed of the Configuration Advisor provide compelling arguments for deployment, and as part of both the autonomic computing initiative the Configuration Advisor demonstrates some of the benefits that autonomic computing promises to bring. All these factors combine to make the Configuration Advisor a valuable tool for database administrators.

References

Eva Kwan, Sam Lightstone, Adam Storm and, Leanne Wu, IBM Server Group, "Automatic Configuration for IBM DB2 Universal Database", online,
<http://www.redbooks.ibm.com/redpapers/pdfs/redp0441.pdf>.

Eva Kwan, Sam Lightstone, Berni Schiefer, Adam Storm, Leanne Wu: Automatic Database Configuration for DB2 Universal Database: Compressing Years of Performance Expertise into Seconds of Execution. BTW 2003: 620-629