

### **Non-Optimized Queries**

**Description:** *The ability to read data from disk, perform some processing on the data, and output the data back to disk, without benefit of caching or optimizations such as primary or clustered indices.*

### **Benchmark Results**

- Cloud-Edge scanned from disk a one (1) billion row, 100-bytes per row, table from disk in 29.619 seconds. The table was conservatively compressed at a 3.7:1 ratio to simulate actual customer data. The benchmark was run on a realistic customer configuration using RAID5 LUNs on a 3Par SAN, connected via standard Brocade 8/24 switches4, resulting in uncompressed throughput of 3.26 Gb/sec.
- Modifying the scanning query described above to include a filter which produced 3.9 million output rows, ran in 29.748 seconds.
- Modifying the scanning query described above to sort the entire table and output it to disk ran in less than 2 minutes (119.167 seconds). The in-memory sort portion of this test was performed in less than 6 seconds.

### **Cloud-Edge Advantage**

The Cloud-Edge database provides a number of built-in features to facilitate the raw processing of large data sets. These include:

- Table compression. Cloud-Edge uses a unique compression technique that typically results in ratios of between 4 and 12 to 1, reducing costs for expensive storage infrastructure and balancing query bottlenecks by taking advantage of available CPUs/cores for on-the-fly compression/decompression.
- Table striping. Cloud-Edge stores a single table across multiple storage logical units (LUNs), allowing for tremendous parallelization of data that is being read or written.
- Massively multi-threaded sort. The Cloud-Edge sort is unique in that it is able to utilize all available CPUs/cores for its highly parallelized sort.

### **Business Advantage**

Analysts performing data mining want the ability to instantly develop and test new ideas. Often these ideas are sophisticated and present extreme processing challenges. Data such as healthcare claims are voluminous and can easily include hundreds of columns, making it impractical to index each of them. Furthermore, aggregating data of this size reduces the level of sophistication that can be developed to discover anomalous patterns of fraud or abuse.

Cloud-Edge on the DL980 presents a unique platform where extremely large data sets can be processed easily and quickly, in whatever manner desired, without the need for tedious data preparation.

### **Description: Legacy application code and data migration**

**Organizations face the challenge of containing costs while delivering new technologies, increasingly complex services, and providing consistently higher levels of customer service. To address these challenges companies create transformation strategies which include the migration of legacy systems, application source code modernization and consolidation, and data migrations to a Pre-Cloud and Cloud-Ready state. The strategies that are created are typically well thought out from a technology blueprint perspective, but often end up being too costly to execute due to lengthy assessments, overall project duration, and the number of resources required to execute the proposed initiative.**

### **Cloud-Edge Advantage**

With Cloud-Edge's modernization platform clients will reduce their modernization initiative costs by utilizing our efficient platform to reduce overall project cost associated with time and resources. With our efficient platform, clients will quickly realize the efficiency of migrating applications from costly platforms like mainframe, dedicated servers, and under-utilized equipment, to a high density Cloud platform in a timely and efficient manner as outlined in our approach below:

The Cloud-Edge platform and methodology takes a unique approach to a reverse engineering process for all legacy objects (i.e. code, data, metadata), and ingests the objects for cataloging and inventory within the Cloud-Edge repository. Once in the Cloud-Edge repository, the advanced analytics engine is utilized to quickly determine the path for the target state application. There are three targeted end-states the Cloud-Edge technology and methodology will drive towards namely:

- Code assessment (Advanced analytical assessment on ingested legacy objects and target application)*
- Code modernization (ex. Cobal, PL1, RPG, Assembly to Java)*
- Data Migration (data structure and definition, data) - legacy structures to relation; legacy data to target*
- Rationalized/Consolidation legacy functions to existing applications including the decommission of legacy modules (ex. UAT/regression testing)*
- Cloud Enablement.*