IBM's IMS — Myths, Realities, and Opportunities

Analyst: Stephen D. Bartlett

Management Summary

There is a significant segment of the population that loves old stuff: old cars and motorcycles, old furniture and fixtures (very popular), and memorabilia of all sorts. These are collected and coveted because of their age, beauty, rarity, emotional connection, or just as examples of how we lived in the past. Personally, I love old architecture; more specifically, older buildings and structures that are unique examples of a style and beauty. You cannot but just love them for what they contribute to the otherwise mundane, cookie-cutter world of the present. There is another set of values that often comes with age – *utility and durability*. An excellent example of this is the U.S. Interstate Highway System, all of those "I" routes that most of us, even if we only do a modest amount of traveling, could not do without.

The system, also known as the *Eisenhower Interstate Highway System* had its origins with the signing of the *Federal Highway Act of 1956*, which created the structure, process, and financing for perhaps the largest public works project in American history. At the time, it was "sold" to the public as a part of our national defense system, enabling the efficient movement of war materiel across the country (and also as emergency runways for aircraft). However, its seemingly unforeseen yet more far-reaching effect was to change the way we live in many ways – the shape and size of our cities – where we worked, played, and vacationed. It created new towns (and killed a few as well) and effectively changed the demographics of the United States. Perhaps the government could have chosen to build more railroads or airports, but roads are far less capital intensive and more flexibly located, to the benefit of us all. Anticipating the challenge for all of the post-Baby Boomers reading this – try to imagine the maps of your town, county, or state with all the Interstate highways erased. Envision how you would travel absent at least a segment of an Interstate highway. *No thanks!* The point here is that regardless of age, utility and durability are values that can be important to our lives.

An example in the IT world –and still much debated intensely and most often fatuously – is the supposed obsolescence of the mainframe and its associated technologies. Submitted as a proof point to allay the concerns of those that still question the validity of IBM's *System z* platform as the core of an enterprise IT infrastructure is the upcoming Golden Anniversary (next year) of its fundamental architecture. Any compiled program written in accordance with the *System/360 Principles of Operation* in 1964 still will execute validly today. To limit the focus and scope of this paper, I will not argue the above case, but assume that the reader is already familiar to some degree with mainframe technology and has a *zEnterprise* system or its predecessors operational in their enterprise or now finds that a zEnterprise system is a valid option for future acquisition. In that domain, however, there is a necessary software complement within which there is a myriad of choices available of vendors and products, the taxonomy of which is too complex to deal with here. In this realm is an area broadly known as *middleware*, because it is software that lies between the operating system and a business-focused application. This set includes database management and transaction management systems that almost certainly are present in one

form or another in any modern IT infrastructure. The "I" in the IBM middleware realm is the *IBM Information Management System (IMS)*. Its age and origins may suggest that this product, or more accurately, the family of products, has long passed its prime and is only sustained by a few hangers-on that can't afford or don't know how to convert their applications to a "more efficient and less costly" platform. Therefore, **the focus of this paper is to disabuse those arguments, and others similar to it, with the express intent of establishing an accurate basis for considering, continuing, or furthering the use of IMS as a viable, if not an essential choice, within your enterprise information processing and delivery strategy.** Please read on to assess this opportunity.

Where does IMS fit in IBM's middleware portfolio?

Within the vast arena of IBM's software portfolio, there is a large set of products categorized as *System z Information Management*. In this group is found the database management, transaction management, and complementary products that assist in the management, controls, and performance of an enterprise's core business functions. This group includes, besides IMS, products such as *IBM DB2 for z/OS*, *DB2 for Linux*, *InfoSphere Information Server*, *InfoSphere Master Data Management*, and *Cognos Business Intelligence for z/OS,* as well as a number of other products associated with information management.

Also commonly included in this set, although not strictly defined by IBM to be in this group, is *IBM CICS Transaction Server for z/OS*[1], also known as *CICS TS*. There is a distinction to be made here, since it is important to the discussion of how IMS fits into an enterprise's information technology strategy. IMS has two complementary product personalities – the database manager (*IMS DB*) *and* the transaction manager (*IMS TM*). These are separate products (i.e., priced and invoiced independently), most frequently used in concert but users are not required to do so. IMS TM is most validly compared to CICS TS, as the essential role of each is the management of messages (transactions) and applications. IMS DB is most validly compared to DB2, as the essential role of each is the organization, storage,

and retrieval of data. The fact is that either of these transaction managers can be the "front end" to either of these database managers, though it would be fair to say that the two IMS products most frequently are paired in concert to provide application solutions for System z users. Together, IMS, CICS, and DB2 are frequently referred to as System z's "legacy" systems because their origins in the case of IMS and CICS go back to the late 1960s, and in the case of DB2, the early 1980s.[2]

What is the best data model for core business applications– hierarchical or relational?

The answer is *"it depends."* Each of these data models is an ideal solution for a particular application framework and workload, and may be rather poor in others. This is what separates the choice between the two. Most readers will be familiar with the relational model because it is ubiquitous, simple to understand, and is supported by a broad range of vendors for almost every server environment in common use today; any programmer worth his/her salt can develop within this paradigm. But the choice is *not* binary, as some relational product vendors would have you believe. I will explain further, later in this report.

The relational data model has data arranged in tables of rows and columns, each row representing an individual record, and each column representing a field or attribute of this record, which are then used as keys to this record. Tables may be logically joined or linked to create a broader or deeper view of the data being accessed. There is a standardized access language, *SQL (Standardized Query Language)*, which allows easy access to relational data. Its strength is in facilitating the retrieval of data for standardized reporting as well as ad hoc query without requiring any knowledge of how the data is specifically organized. As IBM's DB2 product matured, its performance in high-volume, fast-turn-around transactions improved very significantly to the point where it is quite common in these environments to operate in real time without sacrificing its applicability for reporting and query operations.

The *hierarchical data model,* sometimes called a *tree structure*, is much less common in the industry these days and is supported by only a few vendors. Of those products, it would be fair to say that IBM's IMS is – by far – the most widely used and most mature in features and

---

[1] *CICS*, commonly spoken as "kicks", originally was the acronym for Customer Information Control System, which has long since fallen out of use. It became an IBM program product about the same time as IMS.

[2] The origins of these products probably predate the median age of the readers of this paper, which is one of the credibility challenges discussed herein.

functionality. Perhaps the easiest way to think of this model is to visualize a Boeing 787 aircraft, then to imagine how to identify, describe in detail, and keep track of all the pieces, parts, and assemblies and their relationship to one another. This often is called a *bill of materials (BOM)* problem, but in this case it is one involving a very, very large bill of materials. In fact this was the challenge starting in 1966 when IBM partnered with the Rockwell International and the Caterpillar Corporation to develop a data management system best suited to address this problem.[3]

The basic building element of this data model is how the *parent/child relationship* between segments (records) that are linked together. Segments have one or more fields, one of which can be denoted as the *key field* (also called an *index field*) that is used to sequence (order) the record. Each occurrence of a parent segment has associated with it *none* or multiple child segments called *twins*, or *siblings*. A child segment can, at the same time, be a parent segment. The top of the hierarchy, the segment without a parent, is known as the *root segment*. The protocol used to build and access such structures is known as *DL/I (Data Language – One)*. This pragmatic approach can accommodate very complex data structures containing many levels in which one of the primary goals is very fast access, either in a transactional or batch mode of operation, or both.[4] The example used to illustrate the hierarchical data model is typical of those found in product development and manufacturing industries, but has also found extensive use in financial, telecommunications, retail, and cross-industry applications, such as human resources.

So, in answer to the question posed at the beginning of this section, recognizing some of the differences in philosophy between the relational data model and the hierarchical data model, there is ample *room* for both in an enterprise's overall database management strategy. And given their respective strengths and weaknesses, there likely is a *necessity* for both, if the goals of optimized efficiency, lower costs, and responsiveness to the operational needs of the business are to be best

achieved. On these principles, there are innumerable IBM System z clients using both models, as implemented in IBM DB2 and IBM IMS, very successfully over a span of decades.

Present Day IMS Myths and Realities

It seems odd, at least to this observer, that although IBM's IMS and its hierarchical data model predates IBM's DB2 and its relational model by more than a decade, to some IMS may seem to be a forgotten technology. As a consequence, certain stigma seem to be attached to the product that take on the form of "myths" related to obsolescence, complexity, cost, necessary skills, efficiency, and performance. As has been pointed out to me by a respected colleague, one of the advantages of labeling some attribute as a *myth* is that the reader recognizes exactly where the author stands. **If it is a myth, then it is, by definition, untrue, but commonly believed to be true.** Let's begin to dispose of the several myths commonly associated with IMS today.

Myth #1 – The mainframe is not a strategic platform

➢ *Supposed basis for myth* – Continued investments in it or solutions that depend on it are questionable, if not fool-hardy.

➢ *Reasons why this is wrong or exaggerated*
- This myth is at the top of this list to immediately dispose of the notion that the platform on which IMS depends is close to end-of-life. We just need to look at the many global enterprises whose core business operations depend on mainframe technology and who have no plans to leave it.
- Critics often point to a couple of data points that indicate there are fewer mainframes installed today than were installed 10 or 20 years ago. True, but they fail to acknowledge the effects of well-recognized economic behaviors including corporate data center consolidation and the elimination of other redundancies, or the effects of corporate mergers and acquisitions that result in rationalizing and standardizing IT services. In addition, the throughput (execution power) of today's mainframe core is about 150 times more powerful than the first generation CMOS engine. Thus, counting and comparing the number of units of mainframes is a near meaningless exercise. What needs to be measured is the amount of work being done across an

---

[3] The challenge in 1966 was the building of the Apollo 13 spacecraft. It is my impression that Caterpillar's original interest and involvement was for a solution to the large bill of materials processing problems, as it was a manufacturer of many very complex products (and not the space program). See https://en.wikipedia.org/wiki/IBM_Information_Management_System.
[4] It should be pointed out the logical view of data, as seen by the programmer and the user, may bear little or no resemblance to how the data is physically stored and ultimately accessed.

enterprise's mainframes. Comparing MIPS of today's *zEC12* (about 78,000 with 101 user-accessible cores) to the largest first generation CMOS mainframe, (the *System/390 Model 9672*, with 12 cores and 1644 MIPS) tells a different story. The zEC12 can process more than 47 times that workload. Of course, fewer mainframes are needed! And there is no reason that workload performance improvement will not continue.

- From 3Q2010 through 2Q2013, over 210 new System z customers have been added to the fold.
- In the last year, System z installed MIPS growth has been a healthy 23%. Footprint growth has moderated in the developed economies, but the emerging economies of Eastern Europe, China and Southern Asia, and Latin America are experiencing double-digit growth rates. These late adopters have chosen to skip over generations of technology in their eagerness to catch up with the world leaders.

Myth #2 – IMS is not strategic to IBM or to the System z platform

- ➢ *Supposed basis for myth* – There is virtually no continued investment in IMS and the product is being "sunseted".
- ➢ *Reasons why this is wrong or exaggerated*
  - IBM is committed to continuing investment in IMS and its development continues with a new version available every two years.[5] IMS 13 was previewed in October 2012[6], and will be generally available later this year.
  - IMS is considered strategic and essential to the continued success of many enterprise customers around the world; IBM is committed to support these customers.
  - The install base for IMS continues to grow modestly, as new customer licenses are issued every year. The true metrics of its importance are the amount of data being managed by IMS, i.e., using IMS DB, and the number of new applications being enabled to leverage its capabilities, i.e., via IMS

TM. Both continue to grow as enterprise data experiences historic growth rates.

Myth #3 – IMS is old and outdated and not a modern technology

- ➢ *Supposed basis for myth* – It was replaced by DB2 years ago – as all modern database management systems are based on the relational model.
- ➢ *Reasons why this is wrong or exaggerated*
  - This is not substantiated when examining the IMS customer base. As discussed above, the data models have unique strengths that frequently result in having both IMS and DB2 serving the data management requirements of numerous System z customer data centers.
  - IMS TM is often chosen as the front-end for a DB2 database, as it is capable of very high transaction rates and an exceptional quality of service level, particularly in terms of resiliency. It is not unusual for an IMS installation to go for several years without an unplanned outage.
  - IMS DB has been enhanced dramatically, not only to support SQL access in order to integrate well with relational based solutions, but also to offer open access to IMS DB resources directly from z/OS and distributed environments.
  - The IMS Batch Manager is a unique z/OS offering that has built-in checkpoint/restart logic.
  - Not all data management solutions are amenable to the relational model, hence the recent proliferation of *NoSQL* solutions; think of IMS as the enterprise-workhorse NoSQL solution.

Myth #4 – IMS is proprietary and is difficult to integrate

- ➢ *Supposed basis for myth* – It is not "open" and does not "play well" outside mainframe environments.
- ➢ *Reasons why this is wrong or exaggerated*
  - IMS TM, through *IMS Connect*, provides the TCP/IP gateway to IMS transactions, operations, and data that offers the resiliency, scalability, availability, and security essential to enterprise mainframe environments.
  - Simple Object Access Protocol (SOAP), XML web services, JAVA standards, SQL,

---

[5] CICS TS is on an 18-month cycle; DB2 for z/OS has a 3-year cycle.
[6] For more information on IMS 13, go to http://www-01.ibm.com/software/data/ims/v13/.

and other open protocols have been enabled for the IMS ecosystem. It is the goal of the IMS development team to lower if not eliminate barriers to access its facilities, while adhering to the qualities of service that is expected in mainframe environments.

- A good example of the point above is the recognition that IMS must be enabled for access by mobile devices of all types. IBM is openly discussing its intent to deliver enhanced IMS support for mobile applications by implementing support for the *WebSphere Application Server (WAS) Liberty Profile* with integrated REST endpoint, which will enable the use of the lightweight data-interchange format, *JSON (JavaSript Object Notation)*. The support has already been announced for CICS TS.
- The use of *Eclipse*-based application development tools and interfaces for open protocols enforce industry standards for IMS developers to follow.

Myth #5 – IMS is not flexible enough to handle emerging datacenter requirements

➢ *Supposed basis for myth* – Modern businesses require IT solutions that allow them to adapt quickly to new demands and growth in an effective and efficient manner and IMS is too old to be considered modern.

➢ *Reasons why this is wrong or exaggerated*
- IMS provides more efficient CPU and storage utilization, while providing the highest performance capabilities to comparable mainframe-based solutions with no sacrifice to the delivery of the highest qualities of service standards.
- The highest performance, most resilient applications in several industries are built on the IMS TM and IMS DB solution framework. No other solutions have been judged to be able to meet these standards effectively at an acceptable cost.
- IMS DB has three different personalities to address different requirements.
  1. *Full Function* is the most common, historically; this allows direct or sequential access, but is limited to 4 GBs for VSAM and 8 GBs for OSAM.[7]
  2. *High-Availability Large Database (HALDB)* – partitioned parallel processing and built-in online reorganization provides continuous availability for databases up to 40 TBs.
  3. *Data Entry Database (DEDB,* a.k.a. *Fast Path)* provides the highest transaction rates and is partitioned and duplexed for availability. Reorganization is online and batch processing access is not permitted.
- Access to IMS Databases has been enabled from a number of environments:
  1. IMS Transaction Manager
  2. CICS Transaction Server
  3. IMS Standalone Batch
  4. IMS Database Manager (DBCTL control region)
  5. DB2 Stored Procedures
  6. WebSphere Application Server
  7. "Universal" connectivity based on JAVA Database Connectivity (JDBC) API and Distributed Relational Database Architecture (DRDA), both open standards.

Myth #6 – IMS requires very large amounts of costly server and storage resources

➢ *Supposed basis for myth* – IMS only makes sense for the "Big Boys", such as big banks and insurance companies with huge data and transaction volumes, since only they would need it.

➢ *Reasons why this is wrong or exaggerated*
- IMS claims to be the most efficient messaging, transaction, database and batch manager commercially available today. It has been demonstrated that IMS consumes less than 100 MIPS per million transactions.[8]
- IMS physical data takes less storage capacity compared to relational database managers for equivalent amount of user data.
- IMS server resources can be partitioned to fully maximize CPU resources in order to deliver the highest efficiency.

---

[7] VSAM = Virtual Storage Access Method, a native z/OS access method. OSAM = Overflow Sequential Access Method, an IMS-specific access method with performance advantages.

[8] There is a long standing European IMS user in the manufacturing sector that runs its IMS applications in fewer than 100 MIPS, less than 7% of single-engine zEC12.

- IMS handles a wide range of workloads. Volumes range from 100 thousand transactions per day up to as much as 400 million transactions per day in real-world examples.
- IMS makes sense for any size customer that needs an optimized, well-integrated, and cost effective solution platform for z/OS.

Myth #7 – IMS is too complex

➢ *Supposed basis for myth* – It requires a lot of very skilled and experienced technicians to design, manage, and operate an IMS system, and these skills are not available or increasingly hard to find.

➢ *Reasons why this is wrong or exaggerated*
- For many applications, as demonstrated by many System z customers across a number of industry sectors, the hierarchical data model is the most natural and logical, and also provides the best performance. Hierarchical models can be implemented with relational database managers but they are much more difficult to implement and maintain.
- Many open standards have been and continue to be integrated into IMS, as discussed above.
- The *IMS Enterprise Suite Explorer* enables IMS application developers and database architects to perform common and essential tasks easily in an end-to-end application development lifecycle. Using an Eclipse-based user interface, it seamlessly integrates with other Eclipse-based solutions. Graphical editors simplify the development and visualization of essential resource definitions.
- The *IMS* C*atalog* stores metadata about IMS program resources and database resources. This comprehensive view of IMS database metadata, fully managed by IMS, allows IMS to participate in solutions that require the exchange of metadata. It is leveraged by *IMS Open Database*, for open data access across platforms, and *IMS Explorer*, for application development simplification.
- Many of the mainframe skills associated with the management and operation of IMS and IBM's other mainframe products are being addressed through continued development and enhancements to z/OS and *z/OSMF* (*z/OS Management Facility*).

- The IBM Academic Initiative for System z has grown to over 1,000 colleges and universities in 67 countries. The Initiative fosters enterprise computing curricula, certification, and job placement for graduates.

Myth #8 – IMS is not well supported outside of IBM

➢ *Supposed basis for myth* – There are very few ISVs that support IMS and their related product lines are very thin and are not kept current with the latest IMS releases.

➢ *Reasons why this is wrong or exaggerated.*
- Vendors support IMS through a portfolio of tooling products, allowing IMS customers choices from a number of competing products. The Business Partner ecosystem that orbits IMS is very active; more than 25 Business Partners are enrolled in the IMS 13 Quality Partnership Program (the IMS beta program). IMS has a robust user-group ecosystem that not only drives requirements, but in many ways operates as a cooperative with a free exchange of information and experiences. Over twenty regional IMS user groups are active. The SHARE international users group, representing over 2000 of IBM's enterprise computing customers, includes an active IMS project as well.[9]

Myth #9 – IMS costs too much

➢ *Supposed basis for myth* – Like all mainframe-based applications, when compared to distributed systems, those on the mainframe are more expensive to procure, develop, and operate without delivering superior capabilities.

➢ *Reasons why this is wrong or exaggerated.*
- IMS is included in the peripheral damage that surrounds the biases against IBM's mainframe systems.
- IBM Eagle Studies (TCO studies) conducted at hundreds of customer sites with a wider range of application scenarios have demonstrated superior, often dramatically lower, total cost of ownership.[10] Only a small number, about 4%, were an exception.

---

[9] See http://www.share.org/p/cm/ld/fid=43.

[10] These studies are initiated upon request by the customer to assist in platform selection and infrastructure strategic planning. They are not "cherry-picked" by IBM for sales purposes.

- IMS is claimed to be the most efficient messaging, transaction, database and batch manager commercially available today, hierarchical or relational. (See Myth #6 above.)
- To address the issue of software cost for IMS, IBM recently announced *IMS Database Value Unit Edition (VUE)*[11], which introduces a one-time-charge price metric (OTC), plus annual subscription and support charges, for "eligible" workloads. Subject to approval of IBM, eligible workloads may include adoption and integration of smartphones, tablets, and other mobile devices that take advantage of IMS open access capabilities; those in which IMS operational data is included in a business analytics solution; and those implementing JAVA-based applications that access IMS data using SQL. These must be new workloads, not transferred or migrated from another z/OS system, with exceptions for workloads migrated from Linux, UNIX, or Microsoft Windows. Users should also expect substantial discounts to z/OS in the LPARs which host IMS DB VUE-eligible workloads.
- When evaluating the cost of a database management solution, the only truly relevant metric is cost per transaction. A recent study in a U.S. based insurance company has measured their IMS cost per transaction at less than 4/10ths of a cent (< $0.004), which tends to improve as scale increases.
- *What is the opportunity cost if your online mission critical business processes are interrupted for an hour, or a day, or several days? What is the cost of implementing a continuously available infrastructure that will provide disaster recovery with virtually no loss of data?* IBM's mainframe systems with IMS are at the heart of many highly reliable and highly available solutions.

Conclusion

It would be a big mistake to consider IMS as an abandoned technology in IBM's extensive mainframe ecosystem. There are very few contexts where IMS DB, in one form or another, or IMS TM are technologically or economically inappropriate. IMS does not compete with relational database management systems, particularly its younger brother, *DB2 for System z*, but in fact can coexist, complement, and contribute to those environments – as has been demonstrated by many current users. It is clear that IMS is still evolving in concert with the zEnterprise system as it evolves to encompass an expanding hybrid computing architecture.[12]

If System z already is the core of your enterprise IT infrastructure and there are efforts continuing to explore modern technologies that will provide the most effective solutions, then IBM's Information Management System should be high in your consideration set. If your enterprise already has implemented either IMS DB or IMS TM, you should be questioning why your investment is not planned for expansion. Effective arguments have been presented herein that should encourage you to do so.

IMS is ready for the next evolution of mainframe enterprise computing. *Are you?*

---

[11] For more information on the IMS Data Value Unit Edition, see http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&supplier=897&letternum=ENUS213-192.

[12] For more detail on the zEnterprise system see **The Clipper Group Navigator** entitled *The IBM zEnterprise EC12 - Bigger, Better, Faster*, dated August 28, 2012, and available at http://www.clipper.com/research/TCG2012019.pdf and **The Clipper Group Navigator** entitled *IBM's zEnterprise BC12 - More of What You Need*, dated July 23. 2013, and available at http://www.clipper.com/research/TCG2013013.pdf

About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twenty-first year, is an independent publishing and consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

➢ *The Clipper Group can be reached at 781-235-0085 and found on the web at* **www.clipper.com.**

About the Author

Stephen D. (Steve) Bartlett is a Senior Contributing Analyst for The Clipper Group. Mr. Bartlett's interests include enterprise solutions including servers, system software, middleware, their underlying technologies, and their optimal deployment for responsiveness to emerging business requirements. In 2010, he joined the Clipper team after over 42 years with the IBM Corporation as an account and program manager in large system sales, product development, strategy, marketing, market research, and finance. During that time, he received several awards for his contributions in innovative market research and contributions to the business. He has a B.S. from Rensselaer Polytechnic Institute, and an M.S. from Union College.

➢ *Reach Steve Bartlett via e-mail at steve.bartlett@clipper.com or at 845-452-4111.*

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group** *Captain's Log*, **The Clipper Group Voyager**, *Clipper Notes*, *The Clipper Group Calculator*, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *"Navigating Information Technology Horizons"*, and *"teraproductivity"* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosures

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.