KINAXIS™

WHITE PAPER

RapidResponse: How is it so fast?

RapidResponse—How is it so fast?

In today's competitive market, many companies struggle with supply chain management. Outsourcing, globalization and product/channel diversification have created fragmented supply chains, causing significant challenges for inventory management and customer service performance.

Against this backdrop, most companies are faced with high levels of demand and supply volatility. Unexpected changes in demand and supply introduce potential problems and a litany of questions. What is the impact of these unexpected changes to the plan? What is the best corrective action? Answering these questions can take a lot of time and resources. Business success often depends on how quickly companies adjust and implement their plans to maintain the demand - supply balance.

Companies frequently reengineer their sales and operations planning process to gain better agility and alignment. This certainly can help; however, improvement is limited by the software employed. How quickly can you generate and evaluate new scenarios using your planning tool/software? And evaluation of each scenario has little meaning unless you can assess the impact (e.g. to margin and revenue). Do you understand the business impact of your decisions before you make them?

Kinaxis has built specialized supply chain software expertise and experience over the years, which has produced unrivaled solution speed and performance breakthroughs.

At Kinaxis, we have spent over 25-years singularly focused on solving supply chain problems, unlike large software vendors, who offer supply chain planning solutions as a 'complementary' part of their vast list of products. The latter strategy places a lot of strain on the company's core (shared) resources as they have to be divided by the competing product priorities. Furthermore, there is considerable effort spent on trying to integrate the various products. This strategy limits specialization and constrains innovation.

In contrast, Kinaxis has built specialized supply chain software expertise and experience over the years, which has produced unrivaled solution speed and performance breakthroughs. (Appendix A quantifies the superior performance of RapidResponse.)

This paper discusses the technology architecture and the analytical capabilities of RapidResponse that enable its users to create and evaluate multiple 'what-if' scenarios in seconds, while other planning tools can only achieve the same in hours or days. You will learn what sets RapidResponse apart from any other solution in the market in terms of its technology approach to solving supply chain problems.

Technology Architecture

In-Memory Database

RapidResponse uses an in-memory database which is a database management system that relies on main memory for computer data storage, as opposed to more common database management systems which employ disk storage. In-memory databases are faster since accessing data in memory provides faster and more predictable performance than disk.

RapidResponse technology extends this advantage with optimized storage and execution strategies that range from high level software design to low level processor cache considerations.

RapidResponse employs patented technology to efficiently store multiple versions of data using only incremental changes (deltas) in input data.

Efficient Versioning Engine

In order to explore the impact and effectiveness of potential solutions to supply chain issues you need to work from a private copy of your production data. RapidResponse employs patented technology to efficiently store multiple versions of data using only incremental changes (deltas) in input data. This provides for access to numerous scenarios without incurring the large storage cost (see Figure 1) and enables the system to create, store, and compare data from many more scenarios. The lower storage requirements also translate to lower access times and improved performance. These advantages get amplified with more scenarios as multiple users simultaneously explore different issues independently.

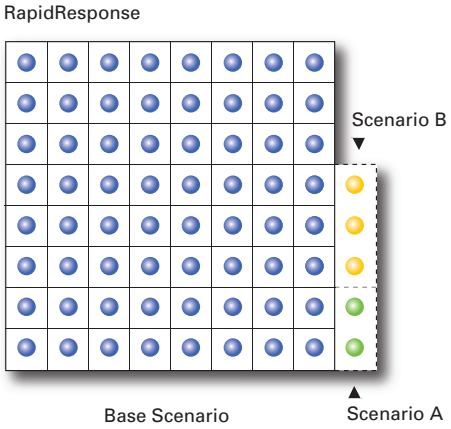

Figure 1: *Storage requirements for scenarios in RapidResponse*

Competitive systems store complete copies of each set of input data. However, the storage requirements to save complete copies of each dataset (or version) can be very large, as shown in Figure 2.

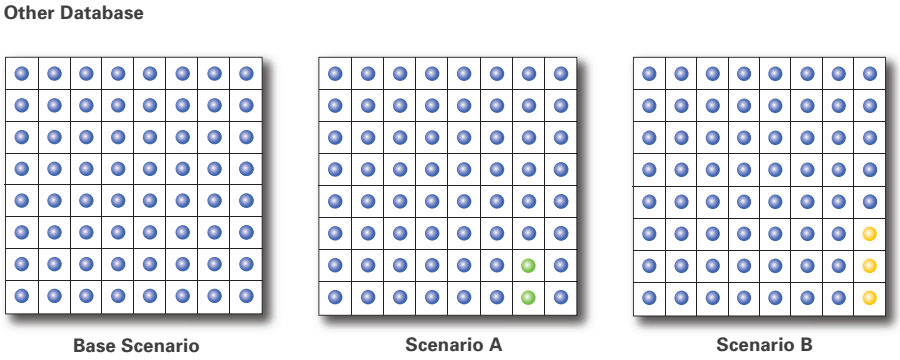

Figure 2: *Storage requirements for scenarios in other databases*

When dealing with complex data relationships like BOMs, the database design employed by RapidResponse provides a significant performance advantage over other competing designs.

Database Design

When dealing with complex data relationships like the Bill of Materials (BOM), the database design employed by RapidResponse provides a significant performance advantage over other competing designs like relational and columnar databases. We will discuss the strengths and weaknesses of relational, columnar, and RapidResponse databases in subsequent sections. Figure 3 provides a summary of the analysis.

	Relational	Columnar	RapidResponse
Search Speed w/ Few Columns	◔	●	●
Search Speed w/ Many Columns	●	◔	●
Efficient Data Storage	◔	●	●
Performance with Complex Relationships e.g. BOM	◔	◑	●

○ *Poor* ● *Excellent*

Figure 3: *Comparing Database Design*

Relational Database

Relational databases are the most common choice for data warehouse applications. In a relational model, data is organized in tables. The relationship with other tables is possible through primary and foreign key relationships. Certain columns (fields) may be designated as keys. A key will use a data structure called an index to speed up data retrieval operations. However, each index requires storage space and reduces the speed of write operations to the associated table.

The relational model allows for query flexibility. There is also a vast array of analytic and integration tools built for relational databases. However, as data volumes grow and users demand more sophisticated analytical capabilities, the performance of relational database decreases noticeably. Relational databases for analytic-intensive applications are difficult to design, challenging to maintain, and require designers to compromise between optimizing query performance and maximizing query flexibility.

The performance of relational models also breaks down when handling complex relationships like BOMs, which defines a product and its components in a manufacturing system. The components themselves may have components that in turn may have other components and so on. A relational database table of all parts will not express the relationships of the parts to the parts of parts. These relationships express important data. To query a database for a product and all its components should be straightforward. A relational database structure makes the developer's job of answering this simple query unnecessarily complex and difficult. In fact, the more complex the collection of information is, the more levels of hierarchy and cross relationships, the more difficult the query becomes to find the parts that collectively make up a product.

Columnar Database

A columnar database stores its content by column rather than by row. A good way to illustrate the difference is using an example of how the same information in a relational database would be stored in a columnar database. Figure 4 shows a part table in a relational database with part name, associated information ('site' and 'type') and a primary key ('ID'.)

ID	Part Name	Site	Type
1	R-FRAME-01	HQ	MRP
2	WH-0521	HQ	MRP
3	DR-0721	HQ	VMI
4	DR-0723	Germany	VMI
5	FR-0122	Germany	MRP
6	WH-0521	Germany	MRP

Figure 4: *Sample Table in Relational Database*

Figure 5 illustrates how the same information can be represented in a columnar database. Each of the separate blocks in this diagram represents separate storage areas. Indexing is usually performed on larger columns and can vary by implementation. Duplicates within a single column can be automatically removed, and null values would not be recorded since the missing record ID implies a null value. This can greatly reduce the storage requirements where there are many instances of data duplication.

ID	Part Name
1	R-FRAME-01
2	WH-0521
3	DR-0721
4	DR-0723
5	FR-0122
6	WH-0521

ID	Site
1-3	HQ
4-6	Germany

ID	Type
1, 2, 5, 6	MRP
3, 4	VMI

Figure 5: *Sample Table Columnar Database*

There are also performance advantages in certain instances when compared to the relational model. Suppose a user searches for information from a few columns on a part table with many records. For example, a user could search for parts from the part table where site is 'Germany' and type is 'VMI.'

A columnar database would search the only two columns ('site' and 'type') of aggregated data for instances where the search conditions are met. Only the row IDs that match all search conditions are returned (e.g. row ID = four). The returned row ID is then used to produce the requested part name. In this case, the performance gain over the relational database is largely driven by spanning fewer columns of aggregated data.

However, relational databases are more efficient when many columns of a single row are required at the same time, and when row-size is relatively small, as the entire row can be retrieved with a single disk seek. So, columnar databases are not always the answer for analytical applications. It is also problematic for columnar databases to handle complex relationships like BOMs as there are time consuming joins employed. There are distinct tradeoffs when considering relational or columnar databases.

RapidResponse uses direct memory references to quickly access the necessary data.

RapidResponse Database

Rather than using time consuming joins that other databases employ, RapidResponse uses direct memory references to quickly access the necessary data. Each reference to a database record in another table is simply a direct pointer to that record. Sets of reverse pointers provide reverse relationship links.

An example of this reference/set relationship is illustrated below. In RapidResponse, each independent demand record for a part includes a direct memory reference to the master data record for that part. Then, the part record for ABC has a set of direct memory references to all associated independent demand records. When a user asks for all independent demand records for a certain part (e.g. ABC), RapidResponse first finds the part and then uses this set to obtain the set of independent demand records. This method for accessing data improves speed significantly. This improvement is even more pronounced with more complex queries and large tables.

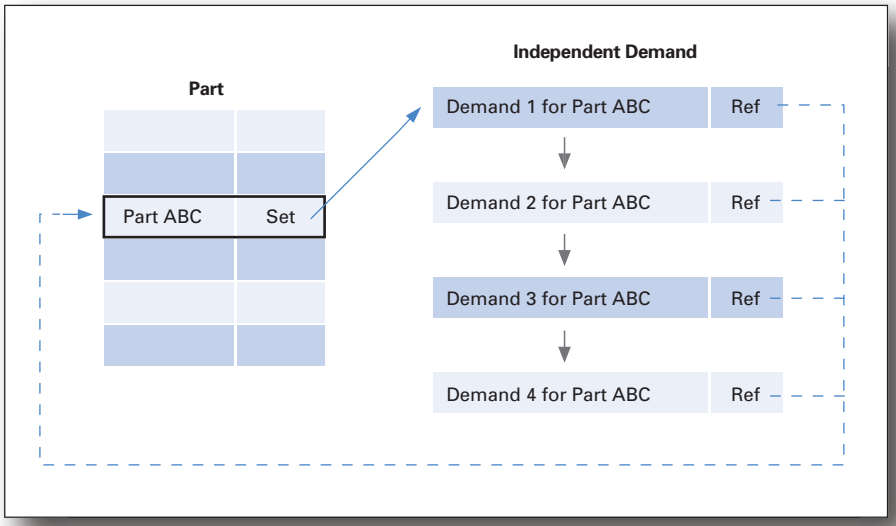

Figure 6: *Reference and Set Relationship*

Optimized Analytics

The analytics code is directly compiled into the RapidResponse database engine where it has direct access to the in memory data and those direct data relationships.

RapidResponse analytics perform highly optimized calculations on the input data and generate calculated output (e.g. planned order recommendations or projected available dates for dropped in orders.) Over the years, we have created and fine-tuned our analytics to take advantage of our high-performance database engine. The analytics code is directly compiled into the database engine where it has direct access to the in memory data and those direct data relationships. Less moving of the data between database and analytics means much better performance. In a relational database this often means copying the data across processes or even across the network. In RapidResponse, calculated output is stored in its own cached data structures, thus avoiding expensive database writes and providing for fast iterative recalculation as one explores different options for input conditions to find acceptable solutions to the supply chain problem that is the focus of the scenario. Invalidated, or least recently used scenario results can be thrown away to make room in memory for the results of currently active scenarios. RapidResponse ships with many analytics.

RapidResponse performs many widely-used analytic calculations that are required for Material Requirement Planning (MRP). This includes the calculation of a part's low-level code, a part's cumulative lead time, netting, explosion, and inventory turns. An example of a useful and widely-used MRP analytic is Multi-Sourcing.

Multi-Sourcing

Multi-sourcing allows a part to have more than one source of supply to satisfy a demand requirement (dependent or independent demand). Multiple sources of supply can include different suppliers, different routings for internal manufacturing (using the same or different product structures), different sites,

and other parts that are equivalent. For example, part ACX-001 can have two sources of supply. The first source may be from supplier ZZZ and the second source may be from supplier RRR. RapidResponse allows for allocation of different sources of supply based on factors such as:

- Target (most commonly used as a percentage)
- Priority
- Contractual requirements

RapidResponse has many unique and proprietary analytics.

When demand is not satisfied by existing supply, a new supply (planned order) is generated (unless the order policies prevent planned orders). When attempting to generate a planned order, RapidResponse determines the source to use based on the sourcing factors. It can even consider shared resource constraints like manufacturing capacity or supplier allocation limits.

RapidResponse also has many unique and proprietary analytics. One such example is Capable-to-Promise.

Capable-to-Promise

The Capable-to-Promise analytic uses a patented method to calculate realistic order completion dates considering component availability, demand priority and capacity throughout the supply chain.

The following calculations are performed by the CTP analytic:

- Supply demand allocation—calculates single level allocations of supply to demand.
- Available date—calculates the date a supply or demand is available based on material availability.
- Gating part—calculates what part, if any, whose lack of availability determines the available date of a demand.
- Incremental availability—splits orders based upon component availability and shows the dates at which these splits could be available (for example, incremental availability might be used to show when different portions of a late customer order will become available.)

Analytic Modifiers

RapidResponse also uses modifiers that can extend and modify analytic calculations through the introduction of additional logic, as well as various infrastructure calculations that are utilized by one or more analytics. Modifiers introduce additional logic to be considered by core RapidResponse analytics, and therefore can impact the output generated by the analytics. For example, the optional Constraint Analysis modifier enables CTP analytics to consider the impact of multiple capacity constraints associated with the sources of supply. Each of the available modifiers can affect some or all of the calculations performed by RapidResponse analytics.

So, how can RapidResponse be so fast?

The technology architecture and optimized analytics provide a strong foundation to achieve the speed and performance demanded by our customers. With each new release of RapidResponse, we leverage our specialized expertise and experience to build on this advantage.

Appendix A – Performance Tests

Test Platform

- Dell Desktop system
- 2.66GHz Core i7 X5650
- 6 cores, 24GB main memory

Test Database

Sample customer data with the following record counts:

Table	*Records*
Part	838,173
Onhand	1,657,205
Scheduled Receipt	521,808
Allocation	1,211,112
Independent Demand	445,623
Bill Of Material	4,055,689
Part Source	1,462,751
	Calculated records
Planned Orders	881,974
Dependent Demand	8,593,897
All supply demand Activity	11,512,214

Results

Raw data access speeds (i.e. parse and execute a query to traverse all the data records and keep a count) on selected table:

Table	Records	Time
Part	838,173	0.037 sec
Bill Of Material	4,055,689	0.170 sec

Simple Query Speeds

- Add up all the quantity fields on the scheduled receipts: **0.050** sec
- Count number of scheduled receipts whose start date is more than 90 "working days" (as defined in the shop calendar tables) before the due date: **0.065 sec**
- Produce a report of the parts sorted by site and name (i.e. a sorted query result with 838,173 records): **7.0 sec**
- Report of the parts summarized by site with a count of the number of parts and the maximum number of scheduled receipts on a single part: **0.741 sec**

Complex Query Speed

- Run complete netting (MRP calculation) and count the number of planned order recommendations: **45 seconds**!
- Running the same calculation again: **5 seconds**! (demonstrates efficient use of cache)

Appendix B – List of Granted Patents and Patents Pending

Patent Granted	Title
1. 7,698,348	Extended database engine providing versioning and embedded analytics
2. 7,610,212	System and method for determining a demand promise date based on a supply available
Patent Pending	**Description**
1. 2009/0171,742	System that schedules higher priority items to complete before target date by handling conflicting requirements

Supply Chain Expert community

Join us online at the Supply Chain Expert Community to LEARN, LAUGH, SHARE and CONNECT. **https://community.kinaxis.com**

ABOUT KINAXIS

Kinaxis delivers a comprehensive on-demand supply chain offering—RapidResponse—that enables manufacturers and brand owners to drive supply chain management (SCM) and sales and operations planning (S&OP) from a single system. Global leaders across a broad range of industries are using RapidResponse as a decision-making hub for the broader value chain and are realizing a competitive advantage as a result. Large manufacturing companies with complex supply chain networks and volatile business environments rely on RapidResponse for collaborative planning, continuous performance management, and coordinated response to plan variances. Learn more about the RapidResponse editions, or join the industry discussions on our Supply Chain Expert community at: http://community.kinaxis.com/.

www.kinaxis.com

Kinaxis World Headquarters
700 Silver Seven Road
Ottawa, Ontario
Canada K2V 1C3
tel: +1 613.592.5780
toll free: +1 866.236.3249
support: +1 866.463.7877
fax: +1 613.592.0584
email: info@kinaxis.com

To the best of our knowledge, this data sheet is accurate as of the date published. This data sheet may be updated by Kinaxis from time to time at its discretion.

2.11