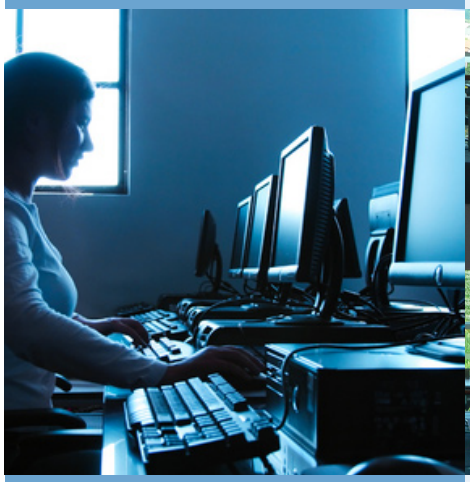

Kinaxis' Self-Healing Supply Chain: Machine Learning in Service of Supply Chain Excellence

Joshua Greenbaum, Principal
Enterprise Applications Consulting
www.eaconsult.com

Spring, 2018

Table of Contents

Introduction: The Journey and the Destination

Getting as close to perfection as possible in supply chain planning and execution has always been the goal, and the computational and modeling power available to planners and the increasing understanding of how to put that power to use has made the push for supply chain excellence more and more within reach.

Nevertheless, there's still an enormous gap between what's theoretically possible and what is being deployed in the real world of supply chain planning. Despite continuous infusions of new technology, a fundamental problem in planning remains: the difference between how the supply chain is designed and how it operates in the real world is often vast. An important starting point is to ensure that plans perform in as realistic a manner as possible, and the key to doing so is to accurately measure actual supply chain capabilities and compare them to the original design. Without the ability to make such a comparison and better ground the plans in reality, plans end up being overly optimistic, unrealistic, or both, and planners are forced to firefight their way through the ensuing mess. This ability to compare actual performance to designed performance also points to opportunities for performance improvements.

This is because the age-old problem of information quality and access still bedevils the world of supply chain planning. Like any complex system, a supply chain's different moving parts are highly interdependent, and a problem with a single component or subsystem – such as a supplier is persistently late, or a part is on world-wide allocation because a new market is consuming the available supply – can have significant repercussions throughout the supply chain. Without visibility into these discrepancies, and without the ability to take corrective action as needed, the supply chain's ability to satisfy demand in a timely fashion diminishes and eventually falters.

In an ideal world, planners would be able to compare their supply chain "as-designed" to its "as-demonstrated" performance – the real world – and make the necessary adjustments to the variables that underlie – adjustments that could serve as a one-time fix or a fix that alters the performance characteristics used to create a plan. In this ideal supply chain world, monitoring a supply chain's behavior and identifying when certain elements of the supply chain are both outside expected performance levels and are significant enough to have an overall impact on revenue, customer satisfaction, overall inventory and other business metrics can help planners make adjustments and reduce the risk. This would help planners to identify the root cause of the discrepancy between the "as-designed" and "as-demonstrated" performance so that corrective action could help return the system to a more ideal state.

This ideal world, it turns out, is closer than it may seem. Recent developments by Kinaxis, Inc. on what the company calls a *Self-Healing Supply Chain* are poised to eliminate this gap between a supply chain's designed performance and its actual performance. Using advances in artificial intelligence and machine learning, a Self-Healing Supply Chain (SHSC) can monitor the different "components" of the system to see if their real-world performance matches the expected performance that is baked into the plan – the "as-designed" state. In doing so, Kinaxis' Self-Healing Supply Chain will be able to flag discrepancies, determine their business impact, and assist planners and others in remediation. Once plans can be

generated based on actual supply chain capabilities, a Self-Healing Supply Chain will eventually be able to remediate specific problems, such as a sudden unexpected spike in demand, the late delivery of a particular order from a supplier, and even macro events such as hurricanes and tsunamis, as they occur.

Kinaxis RapidResponse already provides an important foundational component of this capability by providing a concurrent planning capability. RapidResponse's concurrent planning functionality enables the prediction of the supply chain-wide impact of a discrepancy in the performance of a node in the supply chain. While many large manufacturing companies are already using RapidResponse's concurrent planning capabilities to achieve end-to-end synchronicity in their supply chains, this capability assumes that the design data that is inputted into RapidResponse, usually extracted from an ERP system, is correct. Kinaxis' SHSC initiative begins with the perspective that ERP data may not be correct. By monitoring the supply chain's performance, the SHSC can find the discrepancies in the ERP data that lead the supply chain to deviate from its expected performance. Surfacing these discrepancies and correcting them is the only way we can have confidence that the plan is achievable.

Kinaxis has asked Enterprise Applications Consulting (EAC) to take a first look at its Self-Healing Supply Chain concept and the underlying opportunities and challenges presented by its pioneering work, with a particular emphasis on how customers need to transform in order to realize the self-healing supply chain dream. The first section of this report looks at the opportunity and challenges presented by applying artificial intelligence (AI) and machine learning (ML) to the supply chain planning process. The second section discusses how two Kinaxis customers are already starting down the road to a self-healing supply chain, and what their experience says about how other companies can prepare for this eventuality, and the final section provides some guidance to customers interested in moving their companies forward towards realizing the promise of a Self-Healing Supply Chain.

Self-Healing Supply Chains, Machine Learning, and Data

Kinaxis' work on an SHSC comes at an important inflexion point in the global economy. The complexity and interdependency of global supply chains are growing at a time when the need for greater customer-centricity is driving transformation across all industries. This focus on greater customer-centricity in all its forms is forcing companies to increasingly drive their supply chains towards supporting end-to-end processes that can deliver product where and when it is needed, with the least amount of disruption and the greatest overall efficiency.

This is precisely where the promise of an SHSC will have an enormous impact. Testing the assumptions about lead times and other parameters in the supply chain design against actual historical data will lead to more precise plans that more accurately reflect overall supply chain capabilities. Being able to monitor a supply chain as it runs in real time and make suggestions on how to overcome problems using methods based on a statistical model of a supply chain's design promises to be a vast improvement over ad hoc, more "gut feel" methods. Ultimately, at some point in the near future, a self-healing supply chain will be able to automatically adapt to perturbations in supply and logistics. This semi-autonomous supply chain promises to eke out new efficiencies and improve overall customer satisfaction as never before.

More Than Just New Technology

This focus on customer-centricity in supply chain operations has an important area of commonality with the advances in customer-centricity in ecommerce, retail, and service delivery: a focus on leveraging new and existing data in the service of customers and prospective customers. In this light, considering the quantities of data already percolating throughout the supply chain, it's tempting to look at an SHSC as a simple matter of applying new technology in the form of AI and ML to existing supply chain design, planning and operations. While this simplistic viewpoint is understandable, considering the excess of hype around AI and ML, the prospects for the SHSC need also to be viewed in the context of never-ending, and increasingly important, requirements for accurate and highly curated data.

Digging through the hype, it's clear that all applications of AI and ML are closely intertwined with the quality of the data and information available to the AI/ML system, and that this data quality and quantity issue is the main gating factor for the adoption of these new technologies. At its most basic, an ML system is an intelligent, ever-evolving decision-support system, one that looks at an event or events in real time and then uses a model of "ideal" behavior, built on an AI-based analysis of historical data, to analyze the event and predict a possible outcome or set of outcomes based on the systems' past history. The ML system iterates through this process continuously, and as the system moves into operational mode and more data become available, the ML system is able to use the new data in order to modify or "learn" how to make increasingly better predictions and suggested course corrections for the supply chain.

Users of Kinaxis' RapidResponse are familiar with its modeling functions: Looking through myriad planning scenarios, tracking anomalies, and making course corrections is what RapidResponse does best. The SHSC, by adding the ability to learn and modify suggested outcomes based on the interplay between the "as-designed" state and the real-world state of the supply chain, takes the concept of an intelligent, end-to-end supply chain to the next level.

The graph in Figure 1 shows the output from an SHSC analysis of customer data. The red line is the lead time defined in the supply chain, the "as-defined" state; the individual black dots are the actual lead times, the "as-demonstrated" state; and the blue line is the state of the supply "as-predicted" by the SHSC. The first two-thirds of the data across the x-axis are the historical data that were analyzed by the SHSC, with the resulting blue line being the average deviation from the original design. The last third of the graph is the predicted state of the supply chain moving forward.

This analysis clearly shows that there is a serious discrepancy between how the supply chain was designed and how it actually operates. The designed lead time is five days, but the demonstrated lead time is closer to 30 days. It is clear that using the five-day lead time to determine when to purchase or manufacture new supplies will lead to chronic downstream shortages of materials. The scale of the discrepancy highlighted in Figure 1 could have drastic consequences, including supply shortages and additional costs for expediting supplies, and could ultimately lead to late deliveries to customers, adding costs and putting revenue at risk.

Figure 1: The Supply Chain As-designed, As-demonstrated, and Into the Future

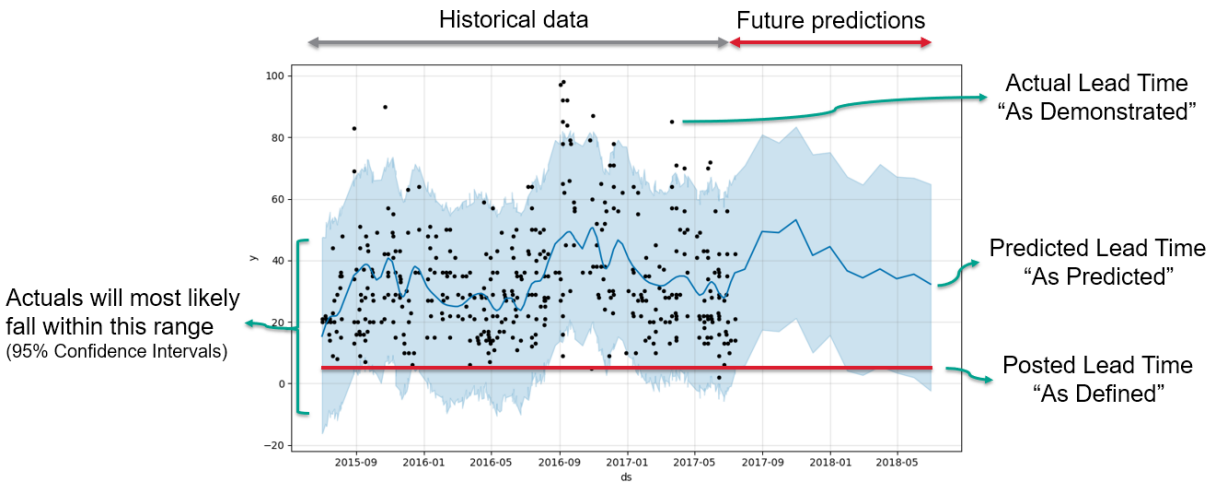

Source: Kinaxis

This initial analysis can be refined even further by looking at the flow of the supply chain using a Sankey diagram (see Figure 2). In the figure below, it becomes clear that the problems with lead times can be traced to a handful of supplies that are chronically late (in red). The width of the line represents the value of the supply, and it's clear that these supplies contribute significantly to overall supply chain performance. From Figure 2 we can clearly see that the biggest issue is the delivery lead time from country ES into country MA, as represented by the red ribbon. The width of the ribbon represents the amount of revenue at risk (in other words, are the supplies likely to be delivered late), and the color represents the degree of discrepancy between the designed and demonstrated lead times, with the shift from green to yellow to red an indicator that lead time is out of spec.

The red ribbon in the chart shows that only a single site (ES) is impacted, the problem is fairly contained, it appears that only one supplier (4400000101) is struggling to deliver on time (red ribbon), and the value of the items delivered is fairly small, as indicated by the narrowness of the ribbon. However, these delays have a big downstream impact on revenue at risk, as indicated by the width of ribbon that flows from ES to MA to AP18.

Figure 2: SHSC Analysis: Lead Times, Key Supplies Down

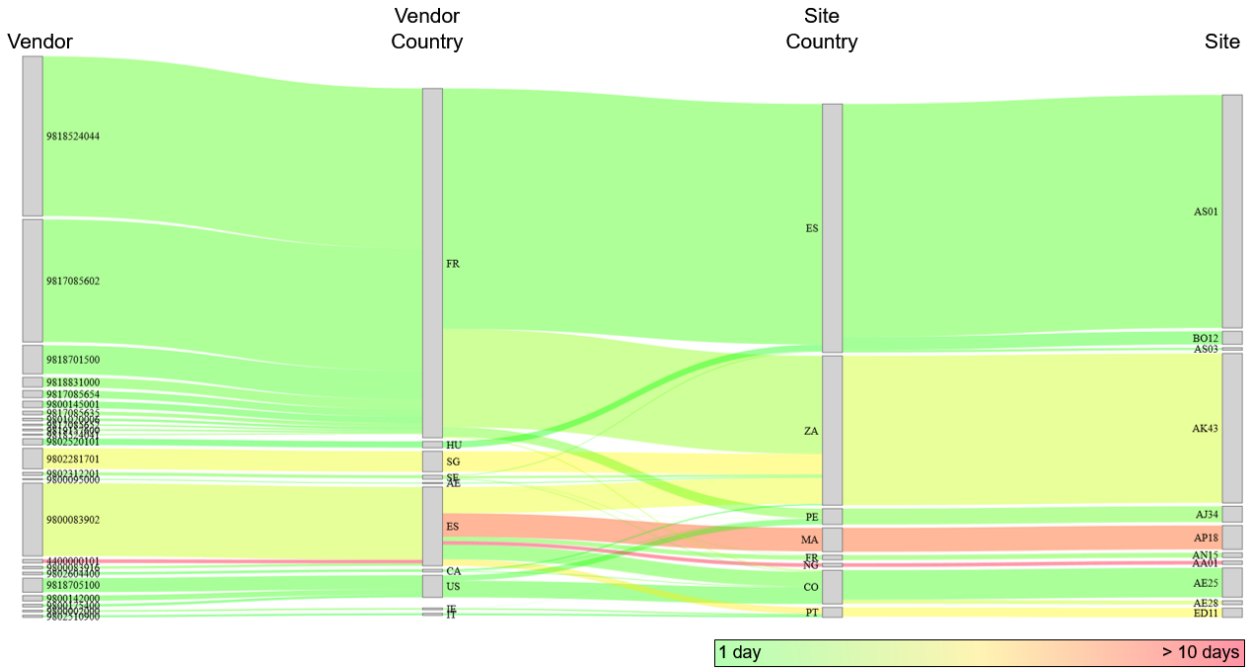

Source: Kinaxis

Evaluating the business impact of the discrepancies is of course a key aspect of the SHSC solution. In the diagram (Figure 3) below, Figure 2 is interpreted in order to highlight the revenue at risk in this supply chain. The pie chart labeled *Part Count* highlights that approximately 30% of the items had a demonstrated lead time close to the designed lead time, meaning that the discrepancies were with an acceptable range (inside tolerance), another approximately 30% of items had greater lead time deviations but where still not in the danger zone (near tolerance), and around 40% of the items were truly outside tolerance.

By itself, this Part Count pie chart doesn't necessarily imply that any action should be taken. That can only be evaluated by understanding business impact, which means looking at the possible interplay between the *Inventory* pie chart and the *Rev @ Risk (revenue at risk)* pie chart. In this case we can see that roughly 80% of inventory is associated with items that are within tolerance, around 10% is associated with items near tolerance, and only 5% of inventory is associated with items whose demonstrated lead time is outside tolerance. Looking just at inventory quantity without an understanding of downstream value, however, doesn't show the full story. The *Rev @ Risk* pie chart shows a much starker picture: Approximately 45% of overall revenue is at risk. This means that somewhere within the relatively small number of parts that are in the red zone in inventory are a set of parts that have a disproportionate impact on revenue. In other words, the analysis shows that remediating lead time problems with a relatively small number of parts can have an enormous return.

Figure 3: Business Impact and Revenue at Risk in an SHSC

Source: Kinaxis

This can be further illustrated by looking at the treemap in Figure 4. While there are multiple dimensions to this map, the key point is that there is an inordinate amount of parts that are late, as all the clusters are red, not yellow or green. But mostly importantly, there is a cluster of parts that are represented by a much darker shade of red. While we can see by its size that this cluster represents a relatively small number of parts, the darker color means that this cluster is not only seriously outside of tolerance, but also places a potentially disproportionate amount of revenue at risk. It's clear from a glance that this darker red cluster is a good place to focus further analysis.

Figure 4: A Treemap View of Lead Time Deviation by Product Cluster

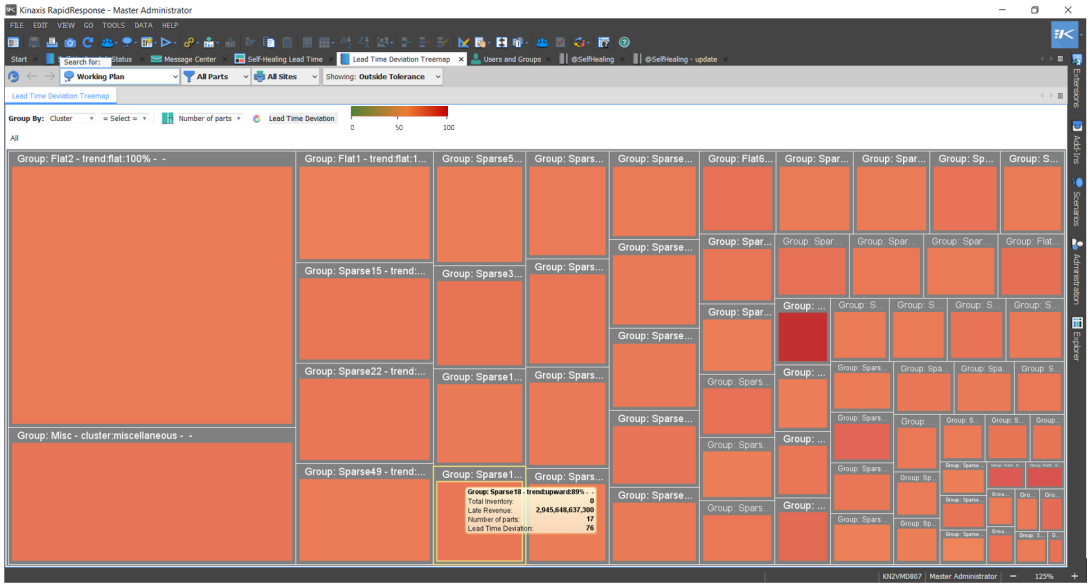

Source: Kinaxis

As we shall see, this is the kind of analysis that customers piloting the use of the SHSC concept are able to generate. EAC believes that these early results, and the remedial action they are able to engender, showcase the enormous potential that an SHSC can have in the real world.

Garbage In, Garbage Out

The promise of the SHSC lies in the improvement of the quality of data used to generate supply chain plans, and the subsequent improvement in the quality of these plans. This opportunity should be of no surprise to seasoned supply chain professionals. Data quality issues have stalked supply chain planning and execution since the dawn of time: *garbage in, garbage out* is an apt description of what happens when a supply chain planning system is fed poor-quality data. This data quality problem is a key aspect of the SHSC concept. By using the AI and ML functions of the SHSC to address the issue of data quality, the broader issue of operating a supply chain semi-autonomously through AI/ML can be addressed.

While this data quality issue, like the overall field of AI and ML and its theoretical application to managing complex systems, is decades-old, the good news is that the underlying technology for deploying AI and ML has evolved significantly. An important innovation has been the improvements in the computational power available for the kinds of complex modeling that form the basis for AI and ML. These advances have opened up the field to its application across a broad range of business problems that are both highly complex and time-sensitive.

A further development has been the greater availability of and ease of access to the necessary data, data quality issues notwithstanding. These new quantities of data span the gamut of formats and types,

an important aspect of the need for new forms of data in order to drive customer-centricity further into the supply chain (see Figure 5). As such, the data needed for an SHSC exist both internally and externally to the company. These data can be highly structured, such as ERP data, and highly unstructured, such as IoT device data, weather data, and customer sentiment data, among many others. Pulling all these data together from their myriad sources has itself necessitated some significant technological breakthroughs, and doing so with the quality and scale needed for AI and ML is an important foundational element in developing an SHSC that is itself continually evolving.

Figure 5: Potential Data Sources for a Self-Healing Supply Chain

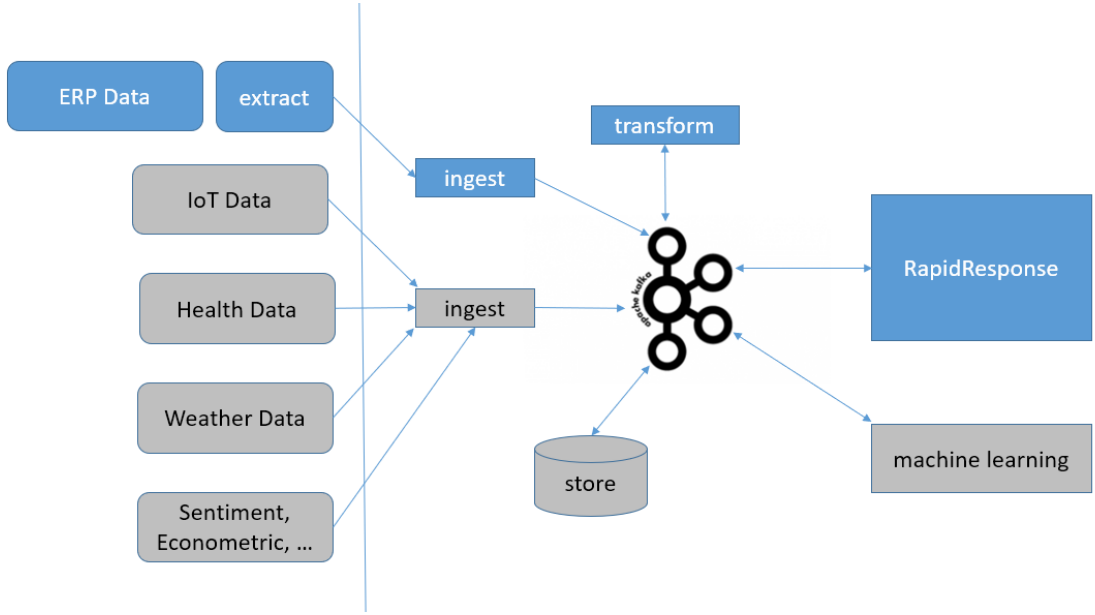

Source: Kinaxis

Presently, early users of SHSC technology are relying on ERP data to build and test their models. As the SHSC concept and technology matures, however, unstructured weather data indicating a possible hurricane along a key supply route, or IoT data indicating a problem in a production facility could be used to yield suggestions for redeploying resources in order meet changes in demand or product supply. Similarly, the analysis of social media data by the SHSC could help fine-tune a planner's understanding of real demand and allow a more intelligent and precise reallocation process than is possible today.

These and other improvements will provide planners in the near-term with suggested course-corrections based on an ML-driven process, and in the long-term will enable an increasingly powerful automated system that takes corrective action without human intervention. While the latter capability is still some years off, it's clear that the advent of the SHSC sets the stage for the ability to "take the robot out of the human" by providing planners with a data and algorithm-based understanding of the impact of the different options available for remediating supply chain problems. Down the road, as experience with an SHSC increases, the more "standard" options will be increasingly automated, leaving planners to focus their attention on problems of greater complexity.

Customers and the SHSC – Initial Results and Findings

Some early customers are already putting the technology and concepts to work, with excellent preliminary results. Kinaxis, through its continuous release cycle, will release extensions and enhancements to meet market expectations.

MSD and the SHSC

For MSD (known as Merck & Co., Inc. in the US and Canada), the prospect of an SHSC fits neatly into the company's plans for a highly customer-centric supply chain. These plans, embodied in a company-wide initiative called Darwin, are intended to dramatically improve MSD's ability "to effectively plan and supply products to our customers at the highest quality, lowest cost and with the shortest lead-time," according to the program's charter.

Complicating these goals is a vast global supply chain with over 15,000 SKUs and over 100 distribution centers servicing customers around the world. With this degree of complexity, the pre-Darwin design and planning processes at MSD, according to Mark Talens, Executive Director Supply Chain Execution, were "designed to fail." "We run into issues that lead to exceptions that lead to firefighting," said Talens. "We don't deal with root causes, we deal with consequences."

The ability to find and deal with the root causes that impact supply chain performance is what drove MSD to use RapidResponse and an early version of Kinaxis' SHSC technology as an important element in its Darwin initiative. "We want to be completely demand driven, and support a more personalized approach to medicine," said Talens. "To do this I am looking more at an end-to-end, proactive supply chain that basically has the capacity to learn and adapt dynamically."

The goal is to have enough intelligence in the supply chain to use the SHSC to understand the real world consequences of a perturbation in the supply chain and adjust the plan accordingly. "If a packaging site misses an order, instead of re-running the order, the system will look forward to seeing what the real demand is," Talens explained. "If the demand isn't there, or if the stock level isn't critical, then the supply chain can keep going with the rhythm."

Beyond Master Data Management

While the goal is a lofty one, the company's SHSC journey, Talens said, started at a relatively mundane level: data. "The single most underestimated element in running a supply chain is data," said Talens. This meant putting in place processes to better manage the quality of the data available to RapidResponse. Importantly, this is much more than just a matter of master data management. Data quality in an SHSC is also about getting the data in line with the intended design of the supply chain and its goals. "You have to look at how the data is set up to support the capital targets of the company," Talens said.

This means that MSD needed to establish a data governance process that looked at how core parameters like lead times or inventory levels are established, and then use RapidResponse and its new SHSC capabilities to start modifying MSD's supply chain design and execution. The SHSC capabilities, in conjunction with RapidResponse, allow MSD to test assumptions about the values of key supply chain parameters based on historical data and then make adjustments based on the statistical probabilities for "success" or "failure" generated by the SHSC.

These initial efforts have already yielded some dramatic results. An effort to take a month's worth of excess supply out of the supply chain is "well underway," Talens reports. Efforts to reduce discards and another effort to reduce "inventory of questionable value" are also yielding quantifiable results, Talens reports.

As the Darwin project rolls out and Kinaxis' SHSC takes on a greater role, MSD is expecting other significant impacts, including the ability to speed up scenario planning, and predict potential stock-outs and other inventory problems. These early successes are helping Talens align the supply chain with the company's aspirations and the goals of the Darwin initiative. "In the end, if we waste less, we serve the patient better," said Talens. "We never forget we're here for the patient."

Schneider Electric and the SHSC

Another early pioneer of the SHSC is Schneider Electric SE, a global energy management automation vendor based in France. Schneider Electric is applying the SHSC technology to its Supply Chain Intelligence Network (SCINET) initiative. As part of that initiative, the company has been running a pilot SHSC program using data from a single region, and the results in terms of visibility into how well – or not – the supply chain functions have been illuminating, according to Brian Tessier, the vice president for Schneider's global supply chain.

The company's analysis of four million records of historical supply chain in its Iberia region was able to identify problems in the supply design that were leading to significant deviations in the supply chain. Lead times were off significantly – in one product group lead times were off by over 50 percent, some of them by 10 days or more. "We could see where it was hitting Schneider," Tessier said at a presentation at Kinaxis' Kinexions user conference last year. "We could come back to the business and say the as-designed lead times are off or you have a service level challenge with a vendor."

Importantly, the SHSC pilot highlighted the fact that the predicted functionality of the supply chain showed "a major disconnect between what we planned and the reality of the situation," Tessier told the audience. "The lead time as predicted is definitely disconnected from what we would expect as a result of how we set this up in the ERP."

The results have spurred Schneider Electric to look at broadening the scope of analysis that can be done by a SHSC. The gating factor, as always, is consolidating data from disparate systems, said Tessier. The pilot was focused on a region that runs a single ERP instance, while many of Schneider's other regional operations are much more complex, with some regions supporting up to 20 separate ERP systems. "The largest barrier we have is the fragmentation of our IT landscape," said Tessier.

Moving forward, Tessier and his team will need to deal with the scale of the company, the scope of its supply, and possible future acquisitions as it continues to roll out the SHSC concept. Dealing with data consolidation is another focus point. Tessier's team is looking at how to consolidate data to make it easier to analyze using non-relational database technologies like HADOOP, and how to ensure buy-in for the SHSC across a global company spanning 44 countries. Buy in is critical, Tessier stated. "After we get all these fantastic findings I have to go back and sell people on the process in order to turn those things into action," Tessier said.

The Self-Healing Supply Chain – Next Steps

Companies interested in following the lead of a Schneider Electric or an MSD can start immediately. There's a lot of work to do in order to realize the potential of an SHSC, and it's not a coincidence that both customers cited above have major business transformation initiatives underway that were well-suited to set the stage for an SHSC. Importantly, the processes involved in setting the stage for an SHSC are necessary pre-conditions for pretty much any major digital transformation, and thus will have payoffs in other important domains as well.

One important basis for digital transformation in supply chain planning, concurrent planning, is already available in Kinaxis RapidResponse. RapidResponse's ability to perform end-to-end supply chain synchronization across data, process, and people/functions in real time is an important first step in deploying an SHSC. Ensuring that RapidResponse has the right data to do concurrent planning – a key starting point on the road to an SHSC – will give companies the on-ramp to develop the skills and processes needed to take full advantage of the ultimate vision for the SHSC, which is semi-autonomous planning.

As noted above, the starting point for an SHSC is data, and any company looking towards implementing this capability needs to focus much of its effort on data quality and data governance. However, making improvements in two other domains – people and process – are also essential.

Data Synchronization, Normalization and Staging: New Processes Needed

Both the theory and the practice of AI and ML require massive amounts of highly accurate data that can be used to train supply chain models and then operationalize them. This is simultaneously both a quantity and a quality problem: If prototyping an SHSC at Schneider Electric for a single region requires four million or more records, imagine what is needed to do this across dozens of regions, thousands of product lines and suppliers, and millions of parts.

The issue of quality and quantity intersect in an important way: many large and not-so-large companies have multiple ERP systems that, even if they are from a single vendor, can vary dramatically from one another. ERP versions can be different, BOMs can be different, and business logic can be different, which means normalizing these varied data sources for use in the kind of analysis needed for an SHSC becomes a daunting task. Even data from a single, global instance can be full of duplicates, errors or omissions.

The result is that data governance is an important discipline for an SHSC, and a non-trivial one at that. Indeed, the overall issue of data governance will by necessity be a continuous and never-ending task: enterprise software and its underlying data models get upgraded, new entities are acquired, business units spun off, and business logic changes. The companies that will succeed in digital transformation will do so by creating and deploying high quality data despite these complexities.

The data quality issue is further complicated by a data access issue: the kind of data needed to train and run an SHSC is by definition highly strategic and contains business critical information that may fall under myriad and possibly conflicting regulatory constraints and/or security requirements. How this data is extracted, aggregated, and staged is an important problem to be solved.

While a creating a single, massive *data lake* is one of the solutions to the access problem under consideration, not only does RapidResponse already support a harmonized view across multiple ERP systems, but building a data lake may not be practical from a legal or business process standpoint. Another option is to maintain data at the source, federate it, and use it in a highly selective manner in order to circumvent the data lake problem. This solution is also made redundant by RapidResponse and runs into the same legal and business process hurdles.

Solving this data quality problem effectively means creating new data governance processes. Supply chains span not just geographies but also internal business units, data silos and other corporate fiefdoms that need to cooperate with one another in order to make the data available for an SHSC while still being compliant with regulatory, security, and privacy concerns that may be regional or business-unit specific. This means that data governance policies have to be put in place that can facilitate the cooperation needed to both maintain the quality of the data as well as ensure its accessibility without compromising other business imperatives, like security. Importantly, these new processes will need to be in place for virtually any significant digital transformation process, whether it's focused on the supply chain or any other aspect of a company's business processes: data will increasingly become one of the most valuable assets a company has at its disposal, and evolving data quality and governance processes is a business imperative for every company.

Planning for the People Side of an SHSC; Change Management and New Skills

There are two basic components to the people side of SHSC: making existing stakeholders part of the evolving SHSC process, and identifying and finding the right talent to enable the SHSC to reach its full potential.

As noted above, the complexity of the data requirement for an SHSC is a reflection of the complexity of the underlying corporate structure, and the requirement that these entities share data and processes means that the individuals in these business units, on both the IT and the business side, provide their buy-in to the changes.

That's not always as easy as it may seem. In fact, the change management side of convincing stakeholders to both participate in the data collection side and then, even more importantly, participate in the remediation processes that result from running an SHSC may end up being the hardest part of implementing an SHSC.

The data cooperation noted above requires business unit stakeholders to be enthusiastic participants in the process, something that may be difficult in a widely dispersed organization. Some of the pushback may come from organizational dysfunction – a lack of a cooperative culture and a tendency to not focus on the larger corporate good in favor of a specific business unit's progress – and some of it may come from an understandable need to ensure that local considerations such as privacy and security be maintained in the face of a program that by definition will be potentially moving petabytes of data outside the business unit's control.

In order to head off what is often a knee-jerk reaction to this level of process change, it's imperative to evangelize the SHSC early and often, ensuring that there's a clear understanding of the value to both the individual business unit, or entity, as well as to the enterprise at large. This stakeholder outreach requirement is no different than what would be needed in any significant digital transformation, and like

the data governance processes discussed above, will serve multiple purposes in the transformation journey, only one of which will be the transformation of the supply chain.

Finally, most companies will need to round out the people side of the SHSC with some new blood or some upskilling of existing personnel. The complexity of the statistical underpinnings of the SHSC will make it increasingly important for supply chain personnel to understand statistics and data governance at a level that is not necessary today. Whether there will be a specific requirement for programming experience with statistical programming languages like R remains to be seen. In theory, Kinaxis and other vendor proponents of the ML and AI revolution are building tools and solutions that don't require direct programming skills. Considering that the concept of the SHSC is so nascent, however, it will certainly behoove pioneering companies to have serious statistics skills in their talent base, and those with programming experience will probably be even more valuable.

Conclusion: The Journey Begins

The work by MSD and Schneider Electric has proven the initial value of the SHSC concept, and the fact that companies are enthusiastically proceeding with their SHSC plans is the ultimate endorsement of the concept. The practical approach each is taking, and their focus on the real world applications of AI and ML, instead of a pie-in-the-sky, overly hyped "robot masters" view of these technologies, reflect each of these companies' thoughtfulness about how an SHSC can proceed. Needless to say, Kinaxis' efforts have followed a similar pragmatism. The result is a program that sets the stage for early and continued success, even while acknowledging the difficulties, particularly in the realm of data quality, that will be present for the foreseeable future.

It's clear that one way or another, the concept of the SHSC will continue to evolve and progress, and while the fully autonomous supply chain shouldn't be in the crosshairs of companies' plans for the near term, not planning for an SHSC is probably something companies will come to regret as time progresses. As stated above, companies at a minimum should be working towards solving the data, process, and people issues inherent to the realization of SHSC for the simple reason that these changes will be needed for all major digital transformation initiatives. The road to digital transformation is paved with data, and most companies have a long way to go in order to get their data houses in order.

It's important to bear in mind that the results of these early pioneers have only scratched the surface in terms of what an SHSC can do, for the simple reason that these two pilots have focused only on ERP data. When the concept evolves to where other external data can be factored into the model – weather, logistics, customer sentiment data, partner data and the like – the true value of the SHSC will be unleashed. These early efforts by Kinaxis and its customers augur well for the future of the SHSC. The question is not *if* the SHSC will be realized, but *when*.