***OR/MS Today* - October 2001**

Supply Chain Management

Effective Collaboration

Hewlett-Packard takes supply chain management to another level

By Gianpaolo Callioni and Corey Billington

More than 10 years ago, Hewlett-Packard realized its first successes using supply chain management to increase its competitiveness and business performance. Supply chain management moved beyond the traditional manufacturing strategy of optimization by individual facility, and began to consider a more holistic view of the entire order fulfillment process. A variety of new applications came from supply chain management, including design for supply chain management (DFSCM), postponement, product rollover strategies and improved inventory fulfillment methodologies.

More recently, supply chain management moved to a new level with the introduction of a collaborative approach that involved multiple partners. One instance of this is the following collaboration between HP and one of its major retail partners, where HP was able to refine the retailer's inventory replenishment strategies to maximize product availability at retail stores without the risk of excess inventory. By using the right level of granularity to classify product sales by both product and location, we were able to transform an unprofitable situation into a highly successful one within a matter of months.

Electronic products present particular market challenges, combining short product lifecycles, a highly competitive consumer-oriented market and low margins. In the early part of the lifecycle having inventory on hand is critical. If the product is not available at the point of sale, the sale is lost. On the other hand, if there is a surplus of obsolete units at the end of a product lifecycle, manufacturers absorb the cost and lose what little profit margin there is.

Supply chain management demonstrated the benefits of optimizing product design for a supply chain, including postponement of customization and localization in favor of generic products that can be customized according to customer demand. It also highlighted the importance of having the right inventory in the right place at the right time, thus allowing HP to cost-justify changes in the supply chain network to optimize inventory distribution.

Now that supply chain management has had time to mature, HP faces the challenge of convincing its strategic partners to cooperate in a new venture: collaboration across multiple nodes of the supply chain among various partners.

The genesis of the collaborative work described in this article lies in the trend within HP toward increased outsourcing. When HP was a single, consolidated company, it was easier to make critical manufacturing decisions, but outside vendors often offered a better combination of flexibility, cost and risk management.

Today, HP outsources to companies that offer these benefits for many manufacturing and distribution processes. However, HP no longer has complete control over the outsourced manufacturing activities; it can no longer unilaterally determine when to turn on and off the production line, what the supplier can produce, or impose its supply chain management practices. The challenge now lies in coordinating the supply chain when a significant piece of it is owned and controlled by partners who can make independent decisions, which nonetheless still affect the performance of the overall end-to-end supply chain.

The Collaborative Mindset

In a traditional, manufacturing arrangement (i.e., one where the nodes don't see themselves as interconnected), each node-to-node relationship is treated as a standalone entity. Between partners, these atomized relationships have additional functional implications that must be addressed in order for collaboration to be effective.

In such arrangements, the sales/purchasing relationship typically acts as the information conduit through which all the functional interactions occur between the partners. In other words, the functions of a particular organization, which may include IT, finance, technical development, manufacturing and shipping, all tend to get their information through this "choke point" (see Figure 1). On the receiving end, purchasing is the activity which all the other organizational functions are driven to support.

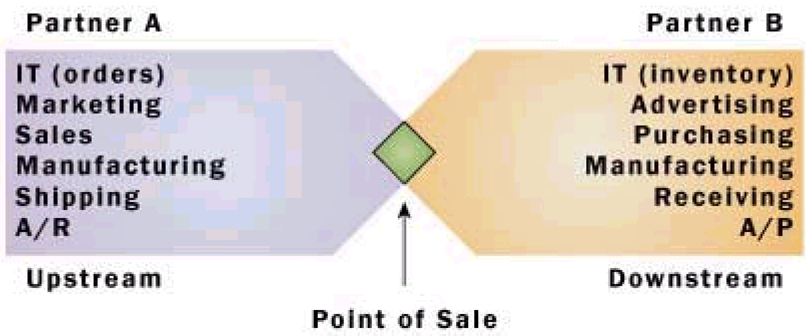

Figure 1: Sales / purchasing as the driver

In the past, supply chain theory has explored the effects of the standalone mindset on activities such as manufacturing, logistics and ordering within the same organization. The first step in supply chain management has been to get all the nodes within the same organization — to see them as interconnected.

However, to create effective collaborations between partners, not only do the organizations need to view themselves as connected, but also the parallel functions within those organizations need to be more closely coupled in a lateral fashion across partners. The supply chain should foster a lateral functional continuity amongst all the partners involved in the collaboration.

This concept was pioneered when Procter & Gamble collaborated with one of their major retail partners, Wal-Mart, to create cross-functional "account teams" composed of individuals from both organizations. Together the team members were responsible for the overall supply chain performance for sales of Procter & Gamble products through the Wal-Mart channel. Figure 2 illustrates how the information channel between the partners widens when complementary functions between the partners are allowed to communicate directly with one another.

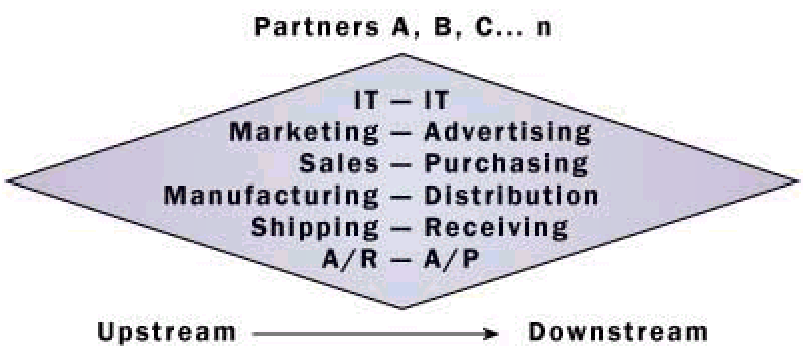

Figure 2: Cross-partner functional integration.

This distinction assumes great importance with regard to information sharing, because it is with the sharing of information that effective collaboration begins.

Methodology

The general methodology for designing a multi-partner supply chain includes these steps:

- Gaining support from all parties by convincing them to view themselves as part of an inter-dependent system.
- Finding the right unit of analysis in order to clearly quantify to each partner the costs of failing to collaborate.
- Implementing the improvements and demonstrating success.

Step 1: Gain Support. The first step is selling the idea. The question to answer for partners is: Why collaborate across a supply chain? Many of HP's supply chains consist of various partners who are best in class at their portion of the supply chain: suppliers, manufacturers, distributors and retailers.

Our earlier projects focused on suppliers and resellers. How could suppliers and resellers be convinced that what is good for one entity is good for all? In particular, how could we win over partners after years of strained, sometime conflicting, relationships driven by the need to contain costs and maximize efficiency in an highly competitive high-tech market? How could we reconcile conflicting metrics?

The academic literature offers several examples of successful partnerships across the supply chain, but only a few articles have investigated the process of developing successful collaboration [Baynes and Shapiro, Corbett and others, et al]. While our findings are quite similar to theirs, our approach was initially aimed to build solid trust through sharing of best practices and through educational efforts. Since our initial interaction, we have continued to share HP knowledge and experience with our selected partners. This has greatly helped in establishing HP as a qualified expert in suggesting alternative processes and methodologies, as well as showing HP to be a trustworthy partner for collaborative projects aimed at improving an entire end-to-end supply chain.

In order to gain their support, we had to persuade our partners to recognize that they were part of an interdependent system. This system took raw goods, transformed them into work in process, and finally into finished goods inventory.

Each partner had to change their viewpoint of themselves from a stand-alone manufacturer, vendor or distributor, and see himself or herself as part of a greater whole. The performance of this whole affected them directly in ways they had not previously foreseen. When they were asked to make changes in the way they did business, each partner would be able to clearly recognize the benefits accrued to them.

This approach had to emphasize the overall systemic benefits as they applied to each member.

The benefits included the following: increased revenue and profit; reduced investment in inventory and other assets; increased product availability (service level) to the final customer; and reduced information distortion across partners.

One pitfall that can occur during this stage is that partners can be caught up in endless discussions about benefit sharing. For example, if one partner shows that a supply chain redesign could result in savings of up to $100 million, the other partners may ask how that $100 million will be divided up, and demand discounts or some other tangible handout that indicates what their slice of the pie will be. If negotiations stall at this point, the collaboration can end before it gets off the ground. To avoid this trap, shift the focus away from shared benefits and on to shared metrics. By using common metrics for measuring success, all partners can measure their own gains as the collaboration proceeds.

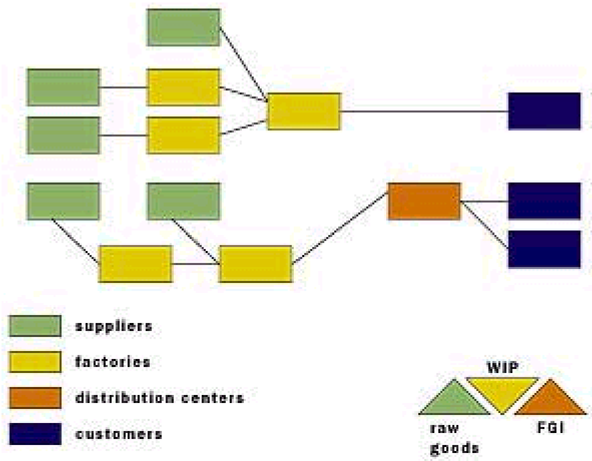

Figure 3: Supply chain as dependent systems.

Depending on the measures chosen, the usefulness of each metric may also need to be explicitly delineated so that the connection between the metric used and the benefits realized by each partner is clearly understood. For example, if revenue is chosen as one of the metrics, the benefits of revenue gains for all partners should be obvious. Increased sales at the retail level translates back upstream into increased orders. However, if product availability is being measured, then the relationship between improved supplier performance and increased sales must be clearly demonstrated to the supplier as well as the retail partners.

Step 2: Finding the Right Unit of Analysis. At this phase, it was essential to get the partners to agree on which shared metrics to use later on when validating new ideas during implementation. Shared metrics are necessary in order to quantify the current costs of a continued failure to collaborate:

- information distortion from order amplification,
- duplication of information at each stage in the supply chain, and
- delays in information pass-through from one supply chain stage to the next.

All of these costs represent opportunities for improvement, through better-shared information between partners: suppliers, vendors and resellers.

Order Amplification. Among the worst of the liabilities is typically the cost associated with order amplification. Order amplification is a phenomenon that occurs when each member of a supply chain "orders up" to buffer his or her own inventory. In the case of a reseller, this means ordering extra stock units. In the case of a manufacturing facility, this means ordering extra raw materials or parts from suppliers, or deploying more expensive capacity. With each stage further upstream, order amplification increases.

Figure 4 illustrates the first phase of order amplification, showing how orders compare to sales for a reseller, in this case for a HP inkjet printer. The amplitude of the sales over time is far smaller than the amplitude or variation in order quantity. In this case, the reseller is ordering extra quantities as a hedge against the uncertainties of delivery, or other factors such as sudden upsurges in demand. The cost of order amplification has been documented in an article titled "The Paralyzing Curse of the Bullwhip Effect in a Supply Chain" [Lee, Padmanabhan and Whang, Stanford University, 1995].

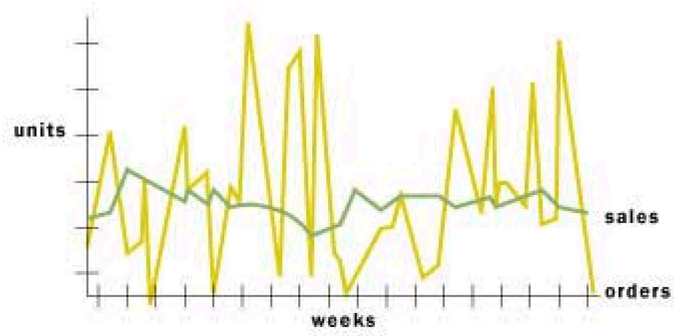

Figure 4: Order amplification

The Bullwhip Effect. Order amplification gets worse the further upstream it is measured. Contributors to this distortion in demand include: independent demand/forecast determinations, time delays caused by order batching, price variations, and guesswork by the downstream nodes regarding supplier performance. Just as a bullwhip's crack has more force at the tip than at the handle, so too does the magnitude of over-ordering increase the farther upstream you go [1].

Figure 5 illustrates this "bullwhip effect." Rather than showing actual units or quantities ordered, the degree of inflation experienced at each stage in the supply chain is shown as a coefficient of variation. This allows us to accurately compare the order amplification of manufacturers, who deliver whole units to resellers, with suppliers, who deal in raw materials. The only measurement of demand that is not an inflation is the sell-through itself. The amplitude of order magnification increases as you go further upstream from resellers through manufacturing, with suppliers experiencing the worst amplification of all.

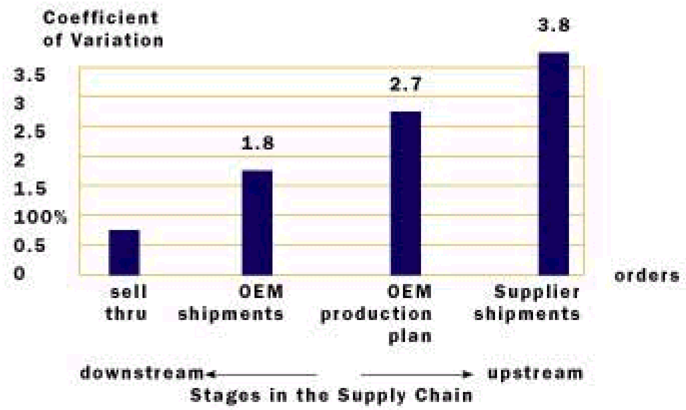

Figure 5: Coefficients of variation increase upstream.

The coefficient of variation is the standard deviation divided by the mean. It serves as a quick way to normalize among different measurements, making them comparable on a meaningful basis. For example, if the mean order is 100 units per week, +/- 20 units, then the coefficient of variation is 20/100, or .2, or 20 percent. The higher the standard deviation compared to the mean, the higher the coefficient of variation; the goal is to lower this coefficient. For example, if you sell 100 units a week and your deviation is 300 — pretty big — then the coefficient of variation is 300/100 or 300 percent. Although this seems unrealistic, this is exactly what is happening in Figure 5: the sell-through coefficient is around 80 percent, the OEM coefficient is around 180 percent, and the supplier's coefficient is 380 percent! This is an order signal amplification almost four times larger than the variability of demand!

Sharing Information. The costs of partners failing to collaborate can thus be traced to lack of access to realistic information, namely, a common demand signal as a driver for upstream supply-chain ordering behavior. Although ideally we would use pure demand sell-through data as the demand signal, in some cases long lead times and short product lifecycles mandate the use of forecasting. For example, a PC with a lifecycle of three months and a part requiring three months to produce requires some forecasting in order to fine-tune production. The proposal here is to have one common pool of information for orders and sales projections, based on proven numbers (sell-through or a unique shared forecast), and to make this information freely accessible to all partners in the supply chain in a form that is most useful to them.

For example, if the reseller and the vendor each maintain separate ordering systems, then every time the reseller places an order, the vendor must reformulate that order and possibly create orders for the supplier as well. This causes delays and duplicates effort. If somehow the sell-through numbers can be translated into the quantities of items most relevant to each supply-chain partner, then each partner can independently adjust their operations according to the same set of information.

Step 3: Make it Happen. The next step is how to put in place a better system. Once you have clearly quantified for all partners the shared benefits to be had from collaboration, and all partners agree on the metrics to use, you should do the following:

- Test the ideas in a pilot project.
- Start tracking your success using the shared metrics agreed upon in Step 1.
- Use results of the pilot to decide what to do for the entire system.

Identifying shared metrics among all partners has proven to be a critical success factor in collaborating with those partners. Instead of endlessly debating how to share benefits and value created by the program, by aligning ourselves around common performance measures we were able to keep the implementation on course and make it happen. The next question becomes: Who should make the critical decisions regarding inventory replenishment, planning and forecasting, and why?

In the following case study, HP took the lead because as OEM it was the most exposed to costs and risks stemming from a failure to collaborate. HP had the best ability to see the entire supply chain and make the best end-to-end decisions.

Retail Case Study

One of HP's retail partners was having a problem with their electronics division, losing sums on the order of $75 million-$100 million a year due to their inability to deploy their inventory to the correct point of sale. Their usually successful supply chain methodology did not take into account the short lifecycle and rapid devaluation of electronic goods. Together, HP and the retailer collaborated through the sharing of data (weekly sales) and the use of a redesigned replenishment process to reduce uncertainty and variability associated with delivery at the retailer's docks.

The real challenge was to increase inventory availability at every point of sale across the entire network. At the same time, we had to contain total inventory levels and leverage the existing network of stores and warehouses.

Their inventory stocking formula was "one to show and two to go," meaning one unit was in the showroom, and two more units were ready to be sold and taken away. This method was too simplistic — the same formula was applied regardless of demand for that particular product at any one location. Thus, some stores were overstocked while others with higher demand never had enough. HP worked with them to set the proper inventory target levels, by store and by product.

The goal of this engagement was to improve the inventory management of the reseller. Without reliable supply and demand information, resellers were forced to order extra stock based on guesswork and loose perceptions of how reliable the suppliers were, and where the demand was likely to be greatest. Despite their best efforts, resellers sometimes found that the inventory was in the wrong place and couldn't be sold, while locations with higher demand were stocked out, thus losing potential sales. With the right information, resellers would be able to order what they needed and direct it where it was needed, instead of ordering more to create unnecessary buffers of inventory.

HP had access to weekly point-of-sale data for each store and delivery performance data for HP as a supplier. Together, these inputs were used to calculate how much inventory to recommend holding at each of the retailer's dozens of warehouses and thousands of retail stores. We performed separate analyses for each product, for several designated points during each product's lifecycle.

HP used internally developed algorithms to optimize the replenishment process. The retailer's stores were classified by the number of units sold to determine how many weeks of supply (WOS) of each product were needed for each class of store. Classification was done separately for each product sold at that location. For example, the same location might be Class A (low volume) for one product and Class C (medium volume) for another.

We were able to prove statistically that mean store demand is a reliable indicator of variability associated with store demand. This relationship allowed us to build our matrix around a simple, well-known number (mean weekly sales) rather than something much more complex such as the variability of weekly sales.

The table shown in Figure 6 is typical of the output that was produced, showing separate store classifications of Low, Low/Med, Med, Med/High and High for each product. Because some stores sell less than one unit per week, having as much as one unit can mean that the "Low" category of stores has a higher WOS number than the "Medium" or "High" stores.

	A	B	C	D	E
Product AA	5.8	2.9	2.1	1.7	1.3
Product BB	4.2	2.4	1.8	1.5	1.3
Product CC	4.6	2.8	2.1	1.5	1.2
Product DD	5.7	2.7	1.8	1.5	1.2
Product EE	...				
...					

Weeks of Supply

Store Classifications

A=Low B=Low/Med C=Medium D=Med/High E=High

Figure 6: Store inventory targets.

The retailer's warehouses helped, because they kept the store inventory targets relatively low, and allowed them to pool the risk demand variability across all stores served by each warehouse. The warehouses also buffered against HP delivery lead-time variability. In our study, we optimized inventory levels for the warehouses as well as each of the stores.

Shared Metrics. Before kicking off any changes in the pilot, both HP and the reseller had to be aligned on the performance metrics. We monitored the following critical metrics to judge the success of this project:

- *Product availability at the stores.* We decided to measure fill rate using product availability. To do this, we used the in-stock rate — checking whether there were any units on the shelf during the weekly stock count (one point-in-time measure per week).
- *Inventory turns in the supply chain.* This measures sales over inventory as weeks of supply. A higher number is better because it means there is less capital tied up in inventory.
- *Product sales.*

The goal was to increase sales while decreasing inventory, because of the costs incurred by maintaining high inventory levels. The higher the inventory, the higher the fill rate, and thus the higher the sell-through (assuming high demand, of course). However, inventory ties up capital. Anything that can reduce the number of months' worth of inventory — while still ensuring that there will be enough on hand to meet the majority of customer demand — is a benefit. Minimizing inventory becomes particularly important later on in a product's life cycle, in order to avoid large inventories of obsolete products.

Inventory Balancing. In this case, product availability within the retailer's network was not enough. The retailer had the right amount of inventory. They were losing money because the inventory wasn't in the right place. Thus, it was important to measure the product availability and sales at each point of sale.

Pilot Results. Figure 7 illustrates a summary of the pilot results for the critical metrics of availability, inventory turns and sales. The measurements took place between May and February of the pilot year, with the pilot beginning in June.

On this graph, the in-stock measurement scale is on the left and the inventory turns is on the right. The sell-through dollar amounts are shown to compare the curve with that for the in-stock and inventory turn trends.

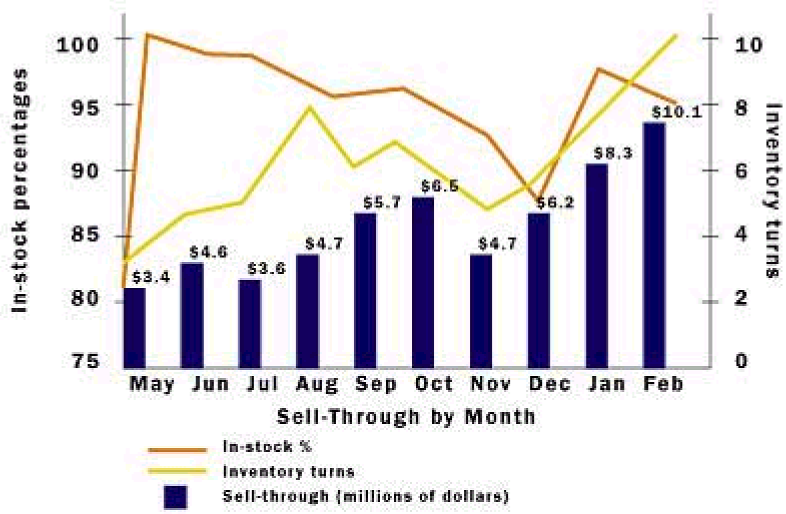

Figure 7: Pilot results.

The results were as follows:

- Before the change was implemented in June, store level in-stock percentage was below 80 percent. It has stayed above 95 percent — the ideal — for 8 of the last 9 months. During the Christmas holidays, we missed our target because there wasn't enough dock capacity to receive HP's products along with the thousands of other products sold in the stores.
- Inventory turns have more than tripled between May and February, from just under 3 to more than 10.
- Sales of these products have increased threefold over the 10-month pilot period, from $3.4 million per month to $10.1 million.

Partner Benefits

Everyone from suppliers and factories, through distribution centers, resellers and customers, stands to benefit from collaboration across the supply chain, leading to performance improvements such as these:

- *Suppliers.* Steadier and more predictable orders means that suppliers are better able to plan their own production schedules, and use any down time for additional work because they don't need to reserve a production line in case of a sudden surge in orders. This reduces uncertainty in capacity planning, which is one of the most costly aspects of suppliers' business operations.
- *Factories.* More reliable deliveries from suppliers and better information on downstream orders mean better production scheduling.
- *Distribution centers and resellers.* More reliable deliveries from factories and fine-tuned inventory fulfillment methods mean they can fill more orders with less inventory, hence lower overhead and fewer lost sales due to unavailability.
- *Customers.* Greater confidence in the suppliers' ability to fill their orders on time or within a known and trusted grace period means better satisfied customers and more repeat business for vendors.

Now that we have successfully conducted a collaborative project across two supply chain partners, we can leverage this approach to envision more far-reaching supply chain collaborations. The ability to synchronize across four or five players, if done correctly, could increase the benefits again for all partners.

Conclusion

The key success factors in using collaborative supply chain management are:

- Move beyond transactional ordering systems to collaborative planning and forecasting for inventory replenishment.
- Look for the right partner company, namely, one that is willing to cooperate in joint planning.
- Agree upon what the critical performance metrics will be, and assure that data is available for those metrics throughout the study.
- Design a modeling system that will meet the needs and capacities of the partners involved, and one that will take into account the nature of the business they are in.
- Design and pilot the processes first, before building automated models.

Measuring Order Amplification

Order amplification can be shown by graphing ordering patterns along with end-user demand (see Figure 4). The metric developed by HP to track this effect [1], called the "order amplification factor," is as follows: order amplification factor equals the standard deviation of orders divided by the standard deviation of demand (sell-through).

This metric compares variability of orders with demand by creating a ratio between the two distributions over the same unit of time. Theoretically the amplification factor can be any value between zero and infinite; however, we expect this number to typically be bigger than 1.

The bigger the order of amplification factor is, the higher the uncertainty experienced upstream; the closer to 1 the order amplification factor is, the more accurately does the reseller's ordering dovetail with actual demand.

Order amplification should be calculated using the same time unit and the same node in the supply chain. For example, compare the weekly sales at node X (a reseller) against orders generated by that same node during the same week. This metric provides a quantitative basis for:

- Comparing customer behavior and ordering amplification for an entire sales channel.
- Quantifying cost of customer ordering policies on supplier inventory and service level.
- Discussion with the customer (reseller) on opportunities to reduce order amplification.
- Measuring improvement over time of initiatives aimed at reducing supply chain costs.

1. G. Gruver, "Supply Chain Inefficiencies in the Computer Industry's Distribution Channels," field research project report, National Technological University, August 1997.

References

1. H.L. Lee, V. Padmanabhan, and S. Whang, "The Bullwhip Effect in Supply Chains," *Sloan Management Review*, Vol. 38, Spring 1997, pp. 93-102.

Gianpaolo Callioni is the director of Strategic Planning and Modeling (SPaM) at Hewlett-Packard. His organization improves the quality of HP's strategic decision-making through the use of quantitative modeling and analysis in areas including supply chain management, forecasting, capacity planning, product design and risk management.

Corey Billington, vice president of Supply Chain Services at Hewlett-Packard, leads a team of more than 1,200 professionals with expertise in procurement, strategic planning, engineering and process management. Previously, Billington was the director of Strategic Planning and Engineering for Supply Chain Services at HP, serving as a key pioneer in HP's approach to supply chain management.

E-mail to the Editorial Department of ***OR/MS Today***: orms@lionhrtpub.com

OR/MS Today Home Page

Lionheart Publishing, Inc.
2555 Cumberland Parkway, Suite 299, Atlanta, GA 30339 USA
Phone: 770-431-0867 | Fax: 770-432-6969
E-mail: lpi@lionhrtpub.com
URL: http://www.lionhrtpub.com