



**NAVAL  
POSTGRADUATE  
SCHOOL**

**MONTEREY, CALIFORNIA**

---

**MBA PROFESSIONAL REPORT**

---

**OPTIMAL DEMAND EXECUTION STRATEGY  
FOR THE DEFENSE LOGISTICS AGENCY**

---

**By: Jeff Bogdanovich and  
Marc LaNore**

**December 2014**

**Advisors: Bryan Hudgens,  
Keebom Kang**

*Approved for public release; distribution is unlimited*

THIS PAGE INTENTIONALLY LEFT BLANK

<b>REPORT DOCUMENTATION PAGE</b>			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
<b>1. AGENCY USE ONLY (Leave blank)</b>		<b>2. REPORT DATE</b> December 2014	<b>3. REPORT TYPE AND DATES COVERED</b> MBA Professional Report	
<b>4. TITLE AND SUBTITLE</b> OPTIMAL DEMAND EXECUTION STRATEGY FOR THE DEFENSE LOGISTICS AGENCY			<b>5. FUNDING NUMBERS</b>	
<b>6. AUTHOR(S)</b> Jeff Bogdanovich and Marc LaNore				
<b>7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)</b> Naval Postgraduate School Monterey, CA 93943-5000			<b>8. PERFORMING ORGANIZATION REPORT NUMBER</b>	
<b>9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES)</b> N/A			<b>10. SPONSORING/MONITORING AGENCY REPORT NUMBER</b>	
<b>11. SUPPLEMENTARY NOTES</b> The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. IRB Protocol number ___N/A___.				
<b>12a. DISTRIBUTION / AVAILABILITY STATEMENT</b> Approved for public release; distribution is unlimited			<b>12b. DISTRIBUTION CODE</b>	
<b>13. ABSTRACT (maximum 200 words)</b>  The Defense Logistics Agency (DLA) updates its demand forecasts and creates orders to replenish aviation hardware inventory levels once each month. The current system creates cyclical shortages in demand planning staff due to monthly spikes in forecasting and purchase order activities. These staffing shortfalls could be reduced or eliminated if the workload were more evenly distributed over time. The project goal is to determine the optimal forecast interval (time between forecast updates) and duration (length of forecast) such that monthly workload variation among the DLA's purchasing staff is minimized, subject to the constraint of satisfying customer order requirements. The measure of effectiveness is the extent to which workload variations are reduced over time compared to the status quo.				
<b>14. SUBJECT TERMS</b> Demand Management, Demand Forecasting, Defense Logistics Agency			<b>15. NUMBER OF PAGES</b> 67	
			<b>16. PRICE CODE</b>	
<b>17. SECURITY CLASSIFICATION OF REPORT</b> Unclassified	<b>18. SECURITY CLASSIFICATION OF THIS PAGE</b> Unclassified	<b>19. SECURITY CLASSIFICATION OF ABSTRACT</b> Unclassified	<b>20. LIMITATION OF ABSTRACT</b> UU	

THIS PAGE INTENTIONALLY LEFT BLANK

**Approved for public release; distribution is unlimited**

**OPTIMAL DEMAND EXECUTION STRATEGY FOR THE DEFENSE  
LOGISTICS AGENCY**

Jeff Bogdanovich, Lieutenant Commander, United States Coast Guard  
Marc LaNore, Lieutenant Commander, United States Coast Guard

Submitted in partial fulfillment of the requirements for the degree of

**MASTER OF BUSINESS ADMINISTRATION**

from the

**NAVAL POSTGRADUATE SCHOOL  
December 2014**

Authors: Jeff Bogdanovich

Marc LaNore

Approved by: Bryan Hudgens

Keebom Kang

William R. Gates, Dean  
Graduate School of Business and Public Policy

THIS PAGE INTENTIONALLY LEFT BLANK

# **OPTIMAL DEMAND EXECUTION STRATEGY FOR THE DEFENSE LOGISTICS AGENCY**

## **ABSTRACT**

The Defense Logistics Agency (DLA) updates its demand forecasts and creates orders to replenish aviation hardware inventory levels once each month. The current system creates cyclical shortages in demand planning staff due to monthly spikes in forecasting and purchase order activities. These staffing shortfalls could be reduced or eliminated if the workload were more evenly distributed over time. The project goal is to determine the optimal forecast interval (time between forecast updates) and duration (length of forecast) such that monthly workload variation among the DLA's purchasing staff is minimized, subject to the constraint of satisfying customer order requirements. The measure of effectiveness is the extent to which workload variations are reduced over time compared to the status quo.

THIS PAGE INTENTIONALLY LEFT BLANK

# TABLE OF CONTENTS

<b>I.</b>	<b>INTRODUCTION.....</b>	<b>1</b>
	<b>A. DEFENSE LOGISTICS AGENCY OVERVIEW .....</b>	<b>1</b>
	<b>B. INVENTORY REDUCTION STRATEGY.....</b>	<b>1</b>
	<b>C. CURRENT FORECASTING AND ORDER EXECUTION .....</b>	<b>2</b>
	<b>D. WORKLOAD VARIABILITY .....</b>	<b>3</b>
	<b>E. SUMMARY .....</b>	<b>4</b>
<b>II.</b>	<b>BACKGROUND AND LITERATURE REVIEW .....</b>	<b>7</b>
	<b>A. BACKGROUND .....</b>	<b>7</b>
	<b>1. Demand Planning at the DLA: An Overview.....</b>	<b>7</b>
	<b>2. Statistical Forecasting at the DLA.....</b>	<b>7</b>
	<b>3. Judgment Forecasting at the DLA .....</b>	<b>11</b>
	<b>4. Summary.....</b>	<b>12</b>
	<b>B. LITERATURE REVIEW .....</b>	<b>12</b>
	<b>1. Demand Management Best Practices.....</b>	<b>12</b>
	<i>a. Demand Sensing .....</i>	<i>12</i>
	<i>b. Demand Shaping and Shifting .....</i>	<i>13</i>
	<i>c. Demand Translation and Orchestration .....</i>	<i>14</i>
	<i>d. Summary.....</i>	<i>15</i>
	<b>2. Inventory Management and Control .....</b>	<b>17</b>
	<i>a. Managing Inventory with Lumpy Demand.....</i>	<i>17</i>
	<i>b. Aggregate Demand Levels .....</i>	<i>19</i>
	<i>c. System-Wide Cost-Oriented Inventory Management .....</i>	<i>19</i>
	<b>3. Forecasting and Staffing Impacts.....</b>	<b>21</b>
	<i>a. Traditional Forecasting Methods .....</i>	<i>21</i>
	<i>b. Non-traditional Forecasting Methods.....</i>	<i>22</i>
	<i>c. Sales and Operations Planning .....</i>	<i>23</i>
	<i>d. DOD Inventory Management .....</i>	<i>24</i>
<b>III.</b>	<b>METHODOLOGY .....</b>	<b>27</b>
	<b>A. BACKGROUND .....</b>	<b>27</b>
	<b>B. ANALYSIS .....</b>	<b>27</b>
	<b>1. Collect DLA Demand Data and Workforce Information .....</b>	<b>27</b>
	<b>2. Evaluate Current Order Execution Strategy .....</b>	<b>27</b>
	<b>3. Generate Hypothetical Workload Models .....</b>	<b>28</b>
	<b>4. Quantify Impacts of Each Model and Compare to Status Quo.....</b>	<b>28</b>
	<b>5. Make Recommendation to the DLA Based on Efficient Staffing..</b>	<b>28</b>
<b>IV.</b>	<b>ANALYSIS .....</b>	<b>29</b>
	<b>A. INTRODUCTION.....</b>	<b>29</b>
	<b>B. CURRENT STRATEGY—THE “STATUS QUO”.....</b>	<b>29</b>
	<b>C. EXPLORING ALTERNATIVE STRATEGIES .....</b>	<b>31</b>
	<b>1. Five-Day Moving Average.....</b>	<b>32</b>
	<b>2. Ten-Day Moving Average .....</b>	<b>33</b>

3.	<b>Thirty-Day Moving Average</b> .....	34
<b>D.</b>	<b>STATUS QUO VS. ALTERNATIVES</b> .....	34
<b>E.</b>	<b>ALTERNATIVE STAFFING REQUIREMENTS</b> .....	40
<b>F.</b>	<b>QUANTIFYING ALTERNATIVES</b> .....	41
<b>G.</b>	<b>IDENTIFYING IMPACTS OF ALTERNATIVES</b> .....	43
<b>H.</b>	<b>SUMMARY</b> .....	44
<b>V.</b>	<b>RECOMMENDATIONS AND CONCLUSIONS</b> .....	45
<b>A.</b>	<b>RECOMMENDATIONS</b> .....	45
<b>B.</b>	<b>CONCLUSIONS</b> .....	46
	<b>LIST OF REFERENCES</b> .....	47
	<b>INITIAL DISTRIBUTION LIST</b> .....	49

## LIST OF FIGURES

Figure 1.	Monthly Purchase Request Generation (R. Wendell, personal communication, July 2014).....	3
Figure 2.	Algorithms used by DLA according to Demand Class (from DLA, 2014).....	9
Figure 3.	Evolution of Demand Management (from Crum & Palmatier, 2003).....	16
Figure 4.	Purchasing and Other Activity vs. Total Time .....	30
Figure 5.	March through June PR Generation.....	31
Figure 6.	Current vs. 5-Day Moving Average.....	32
Figure 7.	Current vs. 10-Day Moving Average.....	33
Figure 8.	Current vs. 30-Day Moving Average.....	34
Figure 9.	Moving Averages Comparison .....	35
Figure 10.	Maximum Daily Workload per Buyer .....	39
Figure 11.	Minimum Daily Workload per Buyer.....	39
Figure 12.	Minimum Buying Staff.....	41
Figure 13.	Annual Work Hours Breakdown .....	42
Figure 14.	Annual Labor Costs .....	43

THIS PAGE INTENTIONALLY LEFT BLANK

**LIST OF TABLES**

Table 1. Input Data for Moving Average Graphs .....36  
Table 2. Sum, Mean, and Standard Deviation .....37  
Table 3. Maximum, Minimum, and Range .....37  
Table 4. Maximum and Minimum Daily Workload per Buyer .....38  
Table 5. Buyer Requirements.....40  
Table 6. Annual Labor Costs and Savings Data .....43

THIS PAGE INTENTIONALLY LEFT BLANK

## LIST OF ACRONYMS AND ABBREVIATIONS

AAC	Acquisition Advice Code (how DLA manages item)
ALT	Administrative Lead-Time (average)
DFU	Demand Forecasting Unit
DLA	Defense Logistics Agency
DOB	Date of Birth (origination date of purchase request)
DOD	Department of Defense
GAO	Government Accountability Office
IT	Information Technology
LP	Linear Programming
PR	Purchase Request
PLT	Production Lead-Time
PTO	Paid Time Off
FSC	Federal Stock/Supply Class
NIIN	National Item Identification Number
S&OP	Sales and Operations Planning
SL	Service Level
SUP_CHAIN:	Supply Chain
STD_U_PRICE	Unit Selling Price DLA charges customers
SWT	Small World Theory
TLT	Total Lead-Time (= ALT + PLT)
TPIP	Time Phased Inventory Plan

THIS PAGE INTENTIONALLY LEFT BLANK

## **ACKNOWLEDGMENTS**

We sincerely thank our spouses and our children for their enduring support during the process of completing our project research and generating this report. We cannot express enough appreciation and thanks for their understanding as we dedicated ourselves to this research and they all sacrificed much.

We would like to thank the major contributors at the Defense Logistics Agency's Office of Operations Research and Resource Analysis (DORRA): Thank you to Mr. Randy Wendell and Mr. Brian Scherrer for never tiring of our endless questions and data calls; and to Mr. Kenneth Mitchell and Mr. Lawrence Vadala for allowing us the opportunity to conduct our research project on this program.

We would also like to thank our advisors, Professors Bryan J. Hudgens and Keebom Kang, for their sage counsel during our many requests for course guidance, and for providing the occasional unsolicited course correction.

Finally, we would like to thank the Acquisition Research Program; especially we thank RADM James B. Greene (Ret.), Dr. Keith Snider, and Ms. Karey Shaffer for supporting our MBA Professional Report.

THIS PAGE INTENTIONALLY LEFT BLANK

## **I. INTRODUCTION**

### **A. DEFENSE LOGISTICS AGENCY OVERVIEW**

The Defense Logistics Agency (DLA) provides combat logistics support, acquisition services and technical services for the Army, Navy, Air Force, Marine Corps, other federal agencies, and allied forces. The DLA is a global enterprise that provides services in 28 countries and is headquartered at Fort Belvoir, VA. The DLA provides almost all of the consumable items that America's military forces need to operate, including 85% of the military's spare parts. In the DLA organization, 25,000 military and civilian employees work together to support 2,400 weapon systems and nine supply chains (DLA, 2013b).

### **B. INVENTORY REDUCTION STRATEGY**

Like many government organizations in recent times, the DLA is facing budget cuts. In August 2013, DLA Director Vice Admiral Mark Harnitchek announced a strategy to cut \$13.1 billion in operating and material costs over the next six years. The DLA cites excess inventory as a major expense within the DLA, estimating that 50% of its \$7 billion inventory is categorized as excess. One way to reduce inventory is to improve the planning and forecast accuracy for purchasing items. In pursuit of this goal, the DLA partnered with LMI, a non-profit consulting firm, to create a software solution that manages DLA purchasing of items with highly variable demand. These highly variable demand items have come to represent a large percentage of DLA inventory because of their unpredictability and because they are inherently difficult to forecast. The new software sets a recommended inventory level that balances the risk of stock out with holding cost expenses and buys to the optimal level instead of chasing a highly dynamic forecast. The LMI solution was implemented in January 2014. In the first four months after implementation, the software reduced inventory levels by 11% for the roughly 800,000 line items with highly variable demand. LMI expects that, ultimately, the software will reduce procurement workload by 50% and save \$180 million in inventory holding costs (Johnson, 2013).

### **C. CURRENT FORECASTING AND ORDER EXECUTION**

Successful forecasting and order execution at the DLA rely heavily on the use of one of several algorithms to accomplish an accurate forecast. This project focuses on the forecasting and order execution strategy for the aviation consumables. The objective is to minimize forecasting and execution workload variability over the course of a given month, while still meeting customer requirements. The DLA employs 708 aviation-specific buyers. Each employee works an average of 2,088 work hours per year, about 61% of which is attributable specifically to buying aviation hardware. The current DLA model forecasts and executes purchase orders at the start of each month with some intermediate purchase orders placed during the remaining days of the month.

The DLA generates PRs every day of the month, but there are spikes that correspond with the times when the DLA updates their forecasts. The updated forecasts can trigger purchase requests (PRs). Demand for some items is so volatile that daily forecasts would generate a PR one day and then cancel that same PR the next day (R. Wendell, personal communication, October 2014). The DLA uses its forecasts to project the re-order point when a PR should be generated. A Time Phased Inventory Plan (TPIP) is computed daily that accounts for the lead-time and current stock on the shelf. Based on the forecasted demand, the TPIP then determines the timing of a purchase request to ensure the stock arrive on-time. Generally, the DLA does not bank or aggregate customer demands; it fills requisitions as soon as they arrive. Usually the DLA does not prioritize its PRs; order execution is largely first come, first served. Occasionally, however, some PRs are held, to smooth the buyer's workload (R. Wendell, personal communication, October 2014).

The DLA has requested this research team to investigate the impacts on monthly workload variability that would result from changing the forecast execution period. The aviation supply chain at the DLA has 24,190 open contracts in FY14, down from 35,963 last year (FY13). Each of these open contracts represents demand for an item that is awaiting contract award. These contracts are segregated into three categories based on dollar amount:

<u>Type</u>	<u>Amount</u>	<u>Award Frequency Goal (per buyer)</u>
Micro	<\$3,000	120 per month
Small	\$3,000–\$150,000	60 per month
Large	>\$150,000	Four per month

(R. Wendell, personal communication, July 2014)

This research focuses solely on the small contracts and the associated purchase requests (PRs). The DLA provided data for four months of PR generation from March through June 2014 in the aviation branch. This data is presented in graphical form in Figure 1 to demonstrate the large workload variability that the DLA experiences every month.

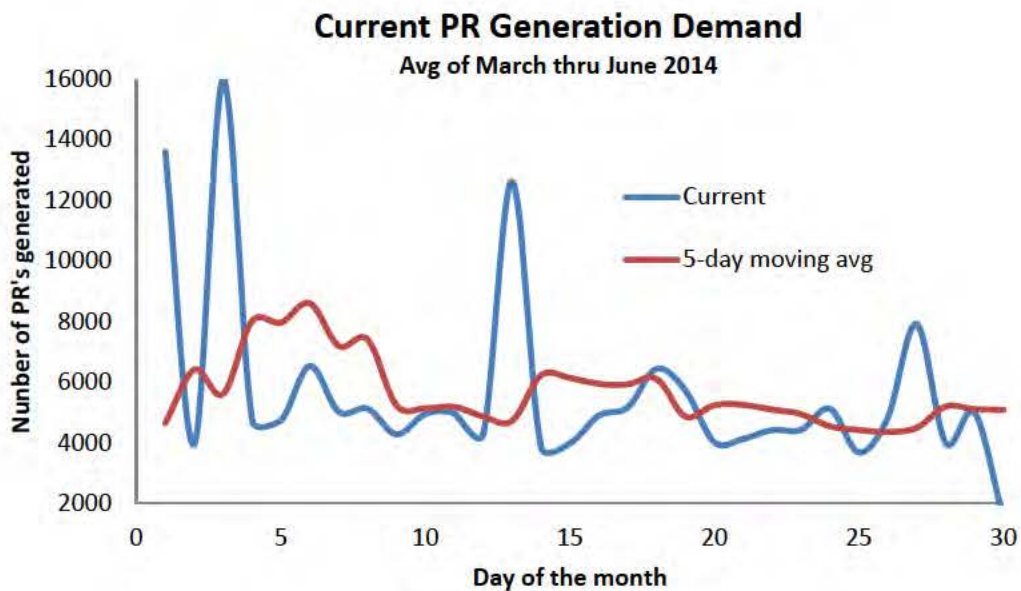


Figure 1. Monthly Purchase Request Generation (R. Wendell, personal communication, July 2014)

#### D. WORKLOAD VARIABILITY

PR processing at the DLA requires a large enough staff to adequately cover the monthly demand peaks. As shown in Figure 1, the staff must be able to handle a daily workload of approximately 16,000 PRs, only to see that workload decrease to about

4,000 PRs a few days later. This huge workload variability is a product of the monthly business cycle in which forecasts are conducted and PRs executed based on those monthly forecasts. The goal of this project is to determine a way to smooth the workload over the course of a month. In Figure 1, a 5-day moving average suggests possible gains from workload smoothing. Assuming that the DLA must staff at the level that accommodates peak demand for any given order execution strategy, then it is easy to see from the graph that staffing requirements for a smoothed workload could be reduced by about 50% of current (708 buyers), or just enough to cover the demand peak on Day 6 of roughly 8,000 PRs. If the workload could be smoothed even more over time, the staffing levels would be further reduced until the highest efficiency point is reached, corresponding to the 30-day moving average.

Workload smoothing is potentially an effective way to save on labor costs associated with PR generation. The key is how to accomplish this workload smoothing. The DLA has requested an investigation into the impacts of adjusting the forecast period and duration to see whether any workload smoothing benefits can be realized. The daily standard deviation was 3,081 PRs for the current business model and 1,146 PRs for a 5-day moving average. In this theoretical example, reducing workload variability by 63% through workload smoothing produced a 50% reduction in labor requirements. There is clear positive correlation between workload smoothing and reduced labor requirements; the challenge is to maximize workload smoothing while maintaining an acceptable customer service level.

## **E. SUMMARY**

In this chapter, we introduced the DLA, provided background on its business model, and emphasized the importance of inventory reduction as a means to lower costs. We also highlighted the importance of forecasting and order execution as they relate to workload variability. In Chapter II, we provide amplifying information about demand management and statistical forecasting at the DLA, as well as demand planning industry best practices. We then adopted a methodology we felt would best review, understand, and analyze the data and information provided by the DLA—outlined in Chapter III. We

develop alternative order execution strategies based upon moving average demand of orders and compare against the current strategy in Chapter IV. We further analyze the various alternatives; we identify quantitative and qualitative impacts of each. In Chapter V, we wrap up the project by providing a recommendation to the DLA on a more efficient workload management strategy that will yield cost savings and reduced variability in order execution.

THIS PAGE INTENTIONALLY LEFT BLANK

## **II. BACKGROUND AND LITERATURE REVIEW**

### **A. BACKGROUND**

In this section, we familiarize the reader with the DLA's current strategy and the important roles that inventory reduction and workload smoothing play in the agency's overall business strategy to cut costs. Moreover, we provide ample information necessary to understand how statistical and judgment forecasting are performed at the DLA.

#### **1. Demand Planning at the DLA: An Overview**

Demand planners at the DLA create a forecast for their supply chain utilizing both judgment and statistical forecasting techniques. The right mix of statistical and judgment forecasting must be applied to develop the most accurate forecast. For demand planners to effectively augment the statistical models with good judgment forecasts, they must understand their products, business, and events. The historical horizon used by the demand planning software is user selectable and includes up to three years of past demand data. The forecast horizon is also selectable by the user and can be varied from two to five years. DLA demand planning takes place on a monthly cycle: Demand history is input once per month, and a forecast is generated and published once per month.

#### **2. Statistical Forecasting at the DLA**

Statistical forecasting uses one of several mathematical calculation algorithms to generate a forecast based on demand history. The mathematical algorithms presume that demand is cyclical in nature, and thus future demand can be predicted by past historical data (DLA, 2013a). The mathematical models and judgment required to make an accurate statistical forecast are complex even when demand forecasting is performed on individual units. The massive volumes of hardware processed by the DLA compound the individual unit analysis and can easily overwhelm a team of even the best demand planners. Consequently, the DLA categorizes its hardware units into one of several pre-defined groups (called demand classes) based on a statistical analysis of historical demand data. Examples of demand classes include continuous, seasonal, continuous non-seasonal, and

erratic non-seasonal. Each demand class has prescribed statistical algorithms and parameter settings that optimize forecast accuracy. Demand planners are responsible for understanding which demand class is most appropriate to the hardware for which they are responsible and applying the mathematical model that is most appropriate for that demand class. This process is called demand classification, and it is the key to creating a realistic statistical forecast.

There are three sub-processes within the demand classification process: classification, optimization, and demand forecasting unit (DFU) creation. In the classification step, DFUs are segmented into one of 10 unique classes based on demand history that includes factors such as seasonality, trend, null orders, and length of history. The following are the 10 classes (DLA, 2013a):

1. Continuous Seasonal
2. Continuous Non-Seasonal
3. Erratic Seasonal
4. Erratic Non-Seasonal
5. Lumpy Seasonal
6. Lumpy Non-Seasonal
7. New Seasonal
8. New Non-Seasonal
9. Obsolete
10. Management Control

After assigning a demand class to a DFU, the optimization process provides a ranking of up to five algorithm recommendations and their associated tuning parameters. Figure 2 shows a table of algorithms by demand class provided by the DLA.

### Algorithms by Demand Class

<b>Continuous Seasonal</b>	<b>Continuous Nonseasonal</b>	<b>Erratic Seasonal</b>	<b>Erratic Nonseasonal</b>	<b>Obsolete</b>
Fourier MLR Holt-Winters (A) Holt-Winters (M) Lewandowski	Lewandowski Holt-SES Fourier MLR	Fourier MLR Holt-Winters (A)	Croston's	Lewandowski
<b>Lumpy Seasonal</b>	<b>Lumpy Nonseasonal</b>	<b>New Seasonal</b>	<b>New Nonseasonal</b>	<b>Management Control</b>
Fourier MLR Holt-Winters (A)	Croston's	Fourier MLR Holt-Winters (A) Holt-Winters (M) Lewandowski	Croston's	User Defined

Figure 2. Algorithms used by DLA according to Demand Class  
(from DLA, 2014)

After DFUs have gone through the classification and optimization processes, demand planners may then map all of the classification and optimization data to the DFU in the creation stage. The creation stage allows the DFU and all of its associated demand attributes to be part of the DLA database.

Once the demand classification process is complete, demand planners can generate statistical forecasts for the newly created DFU with software that automatically selects one of several algorithms. The algorithms currently in use at the DLA include Fourier, Holt-Winters, Croston's, and Lewandowski. The Fourier algorithm works best with stable, seasonal product patterns. It produces good results when two to three years of demand history are used to generate the model and assumes that business changes at a constant rate. The Fourier method fits sine and cosine waves to a time series. The first coefficient in the model ( $C_0$ ) is the mean level of demand. The second coefficient ( $C_1$ ) is the trend, or number of units by which the level is changing each period. Remaining coefficients ( $C_3$ ,  $C_4$ , etc.) represent seasonal patterns in the demand history. The entire

equation is presented in Equation 1, where  $F_t$  is the number of forecasted units demanded as a function of time ( $t$ ).

$$F = C0 + C1 * t + C2 * \sin(\omega t) + C3 * \cos(\omega t) + C4 * \sin(2\omega t) + C5 * \cos(2\omega t) \dots$$

where  $\omega = 2\pi/\text{periodicity}$

The wave with frequency  $\omega$  has one complete cycle per year, while the wave with frequency  $2\omega$  has two complete cycles per year, and so on. (1)

(DLA, 2013a)

Holt-Winters is an exponential smoothing algorithm that is best used to forecast products with continuous, seasonal demand. It uses a weighted moving average to project demand and accounts for level, trend, and seasonality effects in the forecasts it produces.

$$F_t = \alpha X_t + (1 - \alpha) F_{t-1} \quad (2)$$

(DLA, 2013a)

Alpha is the smoothing constant.

The Lewandowski algorithm presumes shifting sales patterns over the life cycle of a product. It calculates a forecast by analyzing actual sales data and detecting patterns in the demand history.

$$X_t = [M_t][S_t][DDE_t][EF_t] + E_t \quad 3)$$

(DLA, 2013a)

$X_t$  is the actual historical value at time period  $t$  and forms the primary basis for the forecast, which is then altered in the formula to account for other effects.  $M_t$  is the value of the dynamic mean.  $S_t$  is the seasonal index. The combination of  $M_t$  and  $S_t$  accounts for seasonality effects in the model.  $DD_t$  is the proportional effect of data driven events, and  $EF_t$  is the proportional effect of external factors. The combination of  $DDE_t$  and  $EF_t$  allows the model to account for time-limited actions taken by a business to influence the forecast. Assumption of  $E_t$  is simply the random error.

Croston's algorithm is used to forecast intermittent and randomly distributed demand patterns with significant null order quantities in the demand history. It is optimal for non-seasonal demand and more than 25% null order quantity.

$$Y_t = X_t Z_t + \varepsilon_t \quad (4)$$

(DLA, 2013a)

where  $Y_t$  is the actual demand occurrence as a function of time  $t$ ,  $Z_t$  is the demand prediction for period  $t$ , and  $\varepsilon_t$  is the random error. Croston's method is a variation of exponential smoothing and is used to predict a mean demand rate when demand is highly irregular. Because the Croston's method only updates the forecast when there is non-zero demand quantity, it does not react strongly to demand occurrences. Unlike traditional exponential smoothing models, it does not trend toward a null value after a series of zero demand states. This characteristic makes it especially suitable for products with intermittent or highly irregular demand.

### **3. Judgment Forecasting at the DLA**

DLA software allows statistical forecasting based on pure historical data where demand planners set the system to calculate the forecast directly from past demand. Since not all demand can be adequately anticipated by a direct mathematical analysis of past events, judgment forecasting is used by demand planners to identify the non-repeatable occurrences in the past and future. For forecasting purposes, a non-repeatable event is an occurrence that will not be included in the base history and statistical forecast. The software has provisions to adjust the forecast based on the demand planner's understanding of this or her products, business, and events. To provide an accurate judgment forecast, the demand planner will identify some portion of the demand history as non-repeatable using a history override and/or create a data-driven event. This approach assumes the event in the past is non-repeatable. The algorithm will generate a statistical forecast from the repeatable history, and a demand planner can add the event into the statistical forecast if identified as a future projected occurrence.

#### **4. Summary**

This section provided the necessary background to understand the status quo of demand forecasting and order execution specific to the DLA. In the subsequent literature review section, we investigate demand planning more broadly in terms of academic understanding and industry best practices.

### **B. LITERATURE REVIEW**

With a thorough understanding of forecasting and order execution at the DLA, the literature review section transitions to a broad examination of demand planning that uses knowledge acquired from industry and academia. Specifically, in the literature review section, we discuss the evolution of demand management, provide examples of several purchasing models, and investigate several best practices from industry leaders.

#### **1. Demand Management Best Practices**

Commercial sector demand management practices are, for many businesses, the fundamental key to success. Cecere (2013) defined demand management as the use of forecasting technologies along with demand sensing, shaping, and translation techniques to improve the supply chain. In the previous section on demand management at the DLA, we examined the forecasting technologies used by the DLA, and now we introduce demand sensing, shaping, and translation.

##### ***a. Demand Sensing***

Cecere (2013) defined *demand sensing* as the translation of downstream data with minimal latency to understand what is being sold, who is buying the product, and how it is impacting demand. By extension, the quality of demand sensing can be assessed by the speed and accuracy of the downstream demand data. Good demand sensing requires the shortest possible time to detect market changes so that the business can respond accordingly with accurate market demand information. Order or shipment data can be used as market sensing tools with a high degree of accuracy, but they are not ideal, as these more traditional data sources have too much latency (Cecere, 2013).

Cecere provided four techniques to improve demand sensing, starting with a focus on the market drivers. In each supply chain, there are one or more market drivers that can serve as leading market indicators. Good demand sensing requires positive identification and tracking of these market drivers. Second, downstream data must be used for successful demand pattern recognition and requires the ability to collect and analyze data across market channels and geographies to understand who is buying which product and in what quantities. Third, after determining the patterns and trends, the demand signal is translated into supply. This can take the form of distribution targets at the distribution center, manufacturing and inventory targets, and supply requirements (Cecere, 2013). Lastly and most importantly in the demand sensing process, demand planners need to close the loop by measuring the impact of demand-shaping programs (Cecere, 2013). This final step determines the extent to which the company has been successful in its demand sensing efforts and serves as a starting point for future modifications.

Successful demand sensing requires the company to quickly and accurately detect and respond to market data. In the next section, we examine how a company might not only react to market demand, but shape the market itself.

***b. Demand Shaping and Shifting***

Good demand sensing relies on speed and accuracy of market data that, combined with the right production response, brings the right products in the right quantities to the market at the right time. It is a highly reactive process that depends on corporate agility for success. Instead of simply reacting to the market, what if the company could shape the market to meet strategic objectives? Not every company can create the exact market for its product, but most can exert some influence over the market, and that is the essence of demand shaping and shifting. Cecere (2013) defined *demand shaping* as stimulating market demand via techniques such as new product launch, price management, assortment, merchandising, sales incentives, and marketing programs.

Demand shifting is similar to demand shaping but it does not stimulate overall demand; it moves existing demand from one period to another. Demand shifting can be a useful tool for a company to align product consumption incentives with factory orders,

distribution, and other logistical considerations. Successful demand shaping relies on three key factors: data analytics, demand pattern recognition, and real-time supply visibility.

Data analytics help companies understand the impact of changing price, trade promotions, marketing events, advertising, and product mix on demand and profitability to make optimal demand shaping decisions (Cecere, 2013). Demand pattern recognition uses customer data and channel insights along with other data to improve decisions (Cecere, 2013). Excellent data analytics and demand pattern recognition are useless unless the company can make rapid and appropriate changes to the supply networks, and that is where real-time supply visibility becomes important. By definition, demand sensing alters demand forecasts, and flexible supply processes translate demand impacts to internal and external supply organizations with minimal signal distortion and latency (Cecere, 2013).

Demand sensing and shaping are analytical tools used to determine the state of the market and manipulate it to meet company objectives. In the next section, we look at some internal controls that allow a company to successfully react to the external market.

### *c. Demand Translation and Orchestration*

We have covered demand sensing and shaping as tools used by a company to understand and control the demand environment. Demand translation and orchestration are analysis and control methods used entirely within the company and focus on the organization's response to the external environment. Cecere (2013) defined *demand translation* as the translation of demand from the market to the organization. The design of this system recognizes that the requirements for demand visibility for each supply chain leader—distribution, manufacturing, and procurement—are different. Successful demand management cannot occur in a vacuum; it is a cross-functional effort that must include open data exchange at all levels, both internal and external to the organization. Each division within a company experiences and must react to market demand fluctuations differently. For example, when demand sensing reveals that product demand is lower, the marketing team must develop incentives that stimulate demand. The

manufacturing division, by contrast, will slow its production in reaction to that same lower demand. Conversely, the manufacturing team must work overtime to provide adequate supply to react to increasing demand, while the marketing team might be directed to quiet its efforts until manufacturing supply can catch up to demand. Thus, different segments within the same organization react differently to the same demand signal, and these efforts must be synchronized and optimized for an efficient outcome; this is the essence of demand translation and orchestration. Cecere (2013) summarized demand orchestration as the successful response to the demand signal that synchronizes the processes of sell, deliver, make, and source across multiple trading partners to determine the right amounts of raw materials to buy and products that need to be shipped.

*d. Summary*

For any organization to perform well in the logistics arena, successful demand management is critically important. Accurate forecasting is a key ingredient of demand management, but it must function as part of an overall collaborative demand management strategy, not a sole solution. As Crum and Palmatier (2003) noted, many companies search for a silver bullet in the form of information technology (IT) as a painless way of forecasting, but IT alone cannot produce accurate forecasts or effective demand management. Successful logistics depends not only on reliable and accurate demand information, but also on how this information is used within the organization to make strategic decisions.

After a reliable demand forecast is performed, the information it provides must be effectively distributed and integrated into business processes at multiple levels within the organizational supply chain, such as procurement, inventory management, and transportation. People play a key role in the process, and there must be continuous investment in their skills and/or education if they are to carry out their pivotal roles effectively. Efficient processes are based fundamentally on collaborative information infrastructure and decision support systems. Crum and Palmatier (2003) reinforced the point that there is no single easy solution to effective demand management; it requires tremendous coordination at almost all levels within the organization.

Demand planning and forecasting was largely unknown 40 years ago. During that time, the prevailing business model was “if you make it, they will buy.” From that early premise, the idea of demand forecasting evolved: The new concept was to attempt to figure out what people want and how much and build only that amount. This was later incorporated at a broader level within the organization and finally, in the last step of the evolution, the demand collaboration was incorporated to include all members of the supply chain. Figure 3 shows the evolution of demand management.

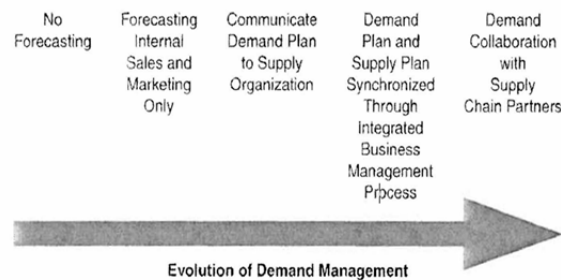


Figure 3. Evolution of Demand Management (from Crum & Palmatier, 2003)

In commercial demand planning, top customers are identified along with the products demanded by these customers that have the highest velocity and earn the highest profits. Given this information, maximizing revenue becomes a simple optimization problem. Part of this optimization model may include carrying reduced inventory for some customers, or discontinuing stocking low demand items. For the DLA, the algorithm is different because of the volume and diversity of product inventory, inconsistent and/or erratic demand, worldwide distribution requirements, and many inventory mandates that are not profitable (R. Wendell, personal communication, July 2014). Given the critical role it plays in the national defense arena, the DLA does not always have the option to simply stop distributing unprofitable parts or refuse shipment to its less valuable customers. The DLA and commercial logistics companies are all motivated to reduce costs, and this is accomplished largely through effective inventory control. In the next section, we discuss the importance of inventory management and introduce several models for accomplishing a cost effective inventory control program.

## **2. Inventory Management and Control**

As demand management is the key to successful revenue generation, inventory management and control play a pivotal role in the cost structure of an organization. There is significant interplay between demand management and inventory management because different demand profiles require different inventory management strategies. Pegels (2005) reinforced this point by stating that market feedback and time lags lead to nonlinearities that add significant complexity to the inventory management picture. In the simple case of steady demand with little variability, the inventory management problem is correspondingly simple with a reorder point at a predetermined inventory level. As demand variability increases, the inventory management problem becomes more difficult. Since steady-state demand drives a trivial inventory stocking question, we focus our attention on the more complicated (and more relevant to the DLA) issue of lumpy demand and how to manage the resulting inventory policy.

### ***a. Managing Inventory with Lumpy Demand***

If inventory management becomes more difficult as variability increases, then lumpy demand represents a significant challenge for inventory managers. Altay and Litteral (2011) summarized that managing inventories when demand is lumpy is complex since companies have to cope with a sporadic pattern that induces high variability in order size that together produce both high inventory levels and unsatisfactory customer service levels at the same time. Further, lumpy demand can cause extreme costs when perceived requirements create excess inventory. According to the Government Accountability Office (GAO), the DLA estimated \$1 billion in inventory for which there had been no demand for at least eight years that resulted in about \$2.5 million in annual storage costs (Edwards, 2011). That same report found that the Army had \$717 million worth of inventory without demand in five years, and that item managers were not performing reviews to eliminate the excess (Edwards, 2011).

The consequences of highly variable demand can clearly be very large in dollar terms if managed improperly. Lumpy demand is the result of the following factors: the

coefficient of variation of demand (CV), the average demand ( $\mu$ ), and the replenishment lead time (LT). Lumpy demand is measured according to the following equation:

$$\text{Lumpiness} = (CV^2) / (\mu * LT) \quad 5)$$

(Altay & Litteral, 2011)

From Equation 5, lumpiness increases linearly with decreases in replenishment lead times or average demands, and lumpiness increases quadratically with increases in demand CV. In addition to this mathematical result, Altay and Litteral (2011) noted several likely characteristics of products with lumpy demand:

- Fewer customers usually induce sporadic requests for the product unit and, therefore, demand lumpiness increases.
- The potential market consists of customers with considerably different sizes or buying behaviors (i.e., customers who order for very different lot sizes or with different frequencies); thus, the higher the heterogeneity of customers, the higher the demand lumpiness.
- There is a low frequency of customer requests. Lower frequency of requests from a customer results in higher numbers of different customers that ask for the unit. Lumpiness increases as the frequency of each customer's purchase decreases.
- There is a high variety of customer requests. Demand lumpiness increases also if each single customer has variable reorder behavior over time.
- There is a high correlation between customer requests. Correlation may be due to imitation and managerial fads, which induce similar behaviors in customers so that sudden peaks of demand occur after periods of no requests. (Altay & Litteral, 2011)

Several of the above lumpiness characteristics can be attributed to product life cycle effects. This is especially true if the parts demanded are limited in quantities. As Altay and Litteral (2011) pointed out, initial spare parts inventories often move slowly off the shelves since requests for them typically fall toward the end of a product life cycle, and requests may change significantly between orders. Early in the product life cycle, the demand for spare parts is low because the parent product is brand new and does not have excessive maintenance needs. As the parent product ages, maintenance requirements increase with correspondingly higher demand for spare parts. As the fleet ages even

more, individual units of the parent product are disposed of, and the demand for spares is reduced. The demand curve over the entire product life cycle approximates an arc with the apex near the middle of the product life cycle. The impact of one parent product life cycle demand fluctuation will not be felt, as spare parts demand lumpiness. Multiple parent products at different stages of the product life cycle with similar sparing requirements can, in the aggregate, create a lumpy demand situation for the common spare parts. We use lumpy demand as the basis to investigate inventory management strategies because spare parts demand often reflects characteristics of lumpy variability.

***b. Aggregate Demand Levels***

It is important to note the importance of the aggregation level of demand. Altay and Litteral (2011) suggested that the demand aggregation level has to be defined over three separate dimensions: the market, the product, and the time horizon. Next, we examine what each of these dimensions means, and then we explain their impacts on demand management and forecasting tasks. Market aggregation here refers to groups of consumers. For example, we can observe the behavior of one consumer (very hard to predict), multiple consumers grouped together by their purchasing behavior at a particular store (somewhat easier to predict), or the behavior of a large mass of consumers that makes up an overall region (least challenging to predict based on the large sample size). Similarly, for product level aggregation definition, it might be very hard to predict the demand for a specific product, while it would be easier to predict the demand for an entire product category given historical demand factors and other existing data. Lastly and most critically important is the forecasting time horizon aggregation of demand; it would be extremely difficult to predict daily demand for a product, but much easier to forecast the yearly demand. Generally, as the aggregation level gets reduced (becomes more specific in product, place, or time), the demand projections become more difficult.

***c. System-Wide Cost-Oriented Inventory Management***

Many companies utilize a Poisson distribution or linear programming (LP) solution to estimate the sparing requirements that then dictate the inventory levels to plug

into the appropriate model. The following are well-known inventory models (Altay & Litteral, 2011):

- continuous review, fixed reorder point ( $r$ ) and order quantity ( $Q$ );
- continuous review, fixed reorder point ( $s$ ) and order-up-to level ( $S$ );
- periodic review, fixed ordering interval ( $T$ ) and order-up-to level ( $R$ ); and
- continuous review, order-up-to level ( $S$ ), and one-for-one replenishment.

Altay and Litteral (2011) stated that the Poisson distribution is far from optimal, as the actual observed stock performance gives low service levels and high costs as a result of over-ordering expensive items following a low-stock event. The priority for management seems to emphasize stock availability, not reducing total inventory cost (Altay & Litteral, 2011). So, the management knee-jerk reaction to a low-stock event actually drives service levels even lower while raising inventory costs for Poisson and LP-based inventory management strategies. Additionally, Poisson and LP solutions can be costly to generate.

Not all companies have the programming capability or the time required to implement accurate solutions for all products and demand profiles within their inventory. For these companies, Altay and Litteral (2011) suggested that companies give a scaled prioritization of inventory holdings prioritized by cost, noting that there was a clear benefit in planning aircraft inventory using the cost-oriented approach. Unlike the Poisson and LP models, which focus on demand requirements and precise spare inventories based on a set service level, the cost-oriented approach focuses on supply costs and prescribes high service levels to low-value parts and low service levels to high-value parts. Altay and Litteral (2011) summarized that by holding more lower-cost items and fewer expensive items, companies can save costs while maintaining overall customer service levels. In this section, we mentioned several useful inventory control models while emphasizing the benefits of a system-wide, cost-oriented approach to inventory management. In the next section, we describe the commonly used forecasting and purchasing models.

### **3. Forecasting and Staffing Impacts**

Inventory managers use many types of forecasting models to assist with demand planning; the range includes complicated analytic tools, judgment of the individual forecasters, simulation models, or some combination of all three that exist within a suitable management framework. In general, the more accurate and specific the analytic forecasting model is, the less likely it will be able to successfully cope with changes to the product, its demand history, or other dynamic events that occur in the real world. Altay and Litteral (2011) echoed this point, stating that the complicated analytical models often have very limited applicability. The tradeoff that exists in developing an accurate forecasting model is thus an exchange between a specific model that applies accurately to a very limited number of cases and a more general model that applies less accurately to a broader range of scenarios.

#### ***a. Traditional Forecasting Methods***

Often the judgment of the human forecaster plays the greatest role in forecasting demand for a particular product. Pegels (2005) wrote that although quantitative forecasting is more accurate than judgmental forecasting, judgmental forecasting is the overwhelming choice of forecasters. The judgment forecast is subject to human errors, but this can be mitigated by the management team through use of the appropriate feedback mechanism. Pegels (2005) explained that there are two main feedback mechanisms in use: (1) performance feedback, which tells the forecaster how well he or she has performed based on forecast accuracy and (2) task feedback, which tells the forecaster how he or she might improve future forecasts. Pegels (2005) stated that if task feedback is provided after each forecast, the accuracy of the forecasts will improve. Additionally, the accuracy of the forecast does not have to rely solely on the quality of the individual forecaster because the judgment of the operators is supported by quantitative measures that reduce the ambiguity of qualitative decisions (Altay & Litteral, 2011).

A combination of two or more forecasting tools can be used to determine a successful forecast. A third choice involves the use of simulation tools that implement the

decisions of the analytical forecast, the judgment forecast, or a combination of the two and run a simulated model to determine the reasonableness of the output. Altay and Litteral (2011) claimed that a simulation model is suitable for the analysis and validation of the decisions made about stock sizing and, when used judiciously, it can improve the outcome. So, it is perhaps a combination of analytical modeling, forecaster judgment, and simulation that gets the best results in forecasting and inventory management. It is also important to note that the management framework plays a role in a successful forecast. Altay and Litteral (2011) explained that a structured decision-making framework can improve overall performance by aligning decisions with a specific business scenario. In summary, a combination of analytical modeling, forecaster judgment, and simulation—all working in synchronicity within a suitable management framework for making decisions—should produce the best results in forecasting and inventory management.

***b. Non-traditional Forecasting Methods***

Most forecasting models rely on demand history to determine an accurate forecast. This works well for products that have extensive, stable histories that can produce a reasonably accurate forecast of future demand. The more difficult products to forecast are those that have short or non-existent demand histories. As Laha and Mandal (2008) pointed out, the main challenge with modeling of demand forecasting for short life span products is the lack of extensive historical data needed to construct an accurate model.

In this environment, models that successfully cope with missing and inadequate data are necessary to construct the gradually degraded model. Laha and Mandal (2008) previewed the use of soft computing techniques such as neural networking to solve the challenging problem of forecasting demand for products with limited demand history. Neural networks are adept at solving pattern classification and prediction (Laha & Mandal, 2008). Neural networking operates on the concept that computers can learn to minimize the errors in data prediction using adaptive algorithm training patterns. Once the neural network has completed a representative learning piece, it can then be released to solve real-world problems that it will also incorporate into its learning algorithm. The

neural network is similar to a human judgment forecaster in that the quality of the forecast is a product of the experience base of the individual performing the forecast.

The neural networks can be utilized to develop models based on the Small World Theory (SWT), which represents the market as a network of interconnected buyers and sellers who may combine to form larger and larger markets as appropriate to the scenario being simulated (Laha & Mandal, 2008). A neural networking model based on SWT is one tool to create potentially accurate forecasts for products with limited demand history. One advantage that SWT has is its capacity to predict demand interplay between competing and complementary products that goes largely ignored by most traditional models. Laha and Mandal (2008) reinforced this point by stating that the local actions of these smaller networks have ripple effects that propagate to the larger market. The development effort involved in creating such a system would be immense—a reasonable undertaking for extremely high-value products with short life spans. After covering the forecasting models, we next discuss supply chain management in commercial and government organizations.

### *c. Sales and Operations Planning*

Many commercial sector organizations utilize sales and operations planning (S&OP). The execution of S&OP involves a mix of inputs from management, sales, operations, finance, and product development. It is a vertically integrated process that connects strategic objectives with master schedules for plants, suppliers, and logistics. There are four elements of a proper S&OP plan: supply, demand, volume, and mix. *Supply* in this context refers to the quantity available to meet the existing and forecast *demand*. *Volume* refers to the aggregate quantities of product families, raw materials, finished goods, and customer backlog. Volume planning is typically conducted on a monthly cycle with provisions for mid-cycle adjustments to accommodate large swings in supply or demand. *Mix* refers to individual products and customer orders. It is more specific and focuses on the individual nuts and bolts, whereas volume involves overall strategy. A successfully executed S&OP plan balances supply and demand while aligning volume and mix in five monthly steps (Wallace, 2006):

- Data gathering and updating: collecting historical results and generating a new statistical forecast
- Demand planning: updating forecasts, adding new product requirements, and creating a consensus forecast acceptable to sales, finance, and production groups
- Supply planning: generating an operations plan that reflects forecast and inventory changes, and identifying capacity problems
- Reconciliation: aligning the demand plan with the supply plan
- Execution: decisions about items unable to be reconciled in earlier phases (Wallace, 2006)

The process is highly cross functional, and each division within the organization must have the right to pose legitimate challenges to the numbers presented by other functional groups if the process is to work effectively. Under S&OP, developing an accurate forecast is a collaborative effort that provides a single product to unite adversarial work groups together with one goal. With S&OP, workload is often reduced because volume forecasts can be extended into the future for 18 months or more while mix forecasts remain more short term and agile (Wallace, 2006).

*d. DOD Inventory Management*

Department of Defense (DOD) supply chain management has often been considered inefficient and ineffective, largely due to questionable forecasting and inventory control practices. Historically, the DOD has been more concerned with avoiding stock out than minimizing inventory cost structure, and that has resulted in huge accumulations of spare parts in excess to needs. A recent GAO report cited inaccurate demand forecasting as the primary reason why DLA accumulates excess inventory (Edwards, 2011). Overestimating demand has led to the excessive inventory that has inflated costs. Underestimating demand has also been damaging. Inaccurate forecasts by the Army and DLA contributed to shortages of parts that caused work stoppages at Army depots in 2006 and 2007 (Edwards, 2011).

It is very challenging for the DLA to forecast the amount of repair parts needed when its depot customers have highly dynamic types and numbers of repairs (Edwards,

2011). Clearly, there are significant costs associated with both over- and underestimating demand, in addition to the resulting high inventory carrying costs. Part of the problem lies within insufficient communications between the DLA and the services, or the customers. The DLA retained large amounts of excess contingency stock, but the services had not provided the input the DLA needed to lawfully eliminate the excess inventory (Edwards, 2011). To resolve these issues, the GAO issued a plan as part of its report whose objective was to improve the demand forecasting such that inventory requirements could be more closely linked to actual needs. To achieve this objective, the DOD identified several actions (Edwards, 2011):

- (1) Identify improved methods and techniques for demand forecasting that consider an item's life cycle (i.e., new item introduction, sustainment, and end-of-life);
- (2) Implement standard metrics to assess forecasting accuracy and bias;
- (3) Expand and refine a DOD-wide structure for collaborative forecasting;
- (4) Implement approaches for improving the setting of inventory levels for low-demand items; and
- (5) Examine how investment risk for new consumable items initially entering the inventory can be reduced among the services, DLA, and suppliers. (Edwards, 2011)

The DLA has examined each of these and is continuing to make significant progress. This project methodology is focused on forecasting execution and its impact on staffing workload variability.

THIS PAGE INTENTIONALLY LEFT BLANK

### **III. METHODOLOGY**

#### **A. BACKGROUND**

The DLA, like many other government entities, is under significant congressional pressure to reduce program costs through efficiency initiatives. As noted in earlier chapters, many functions of the DLA are out of its direct control, and as such, the DLA has no ability to affect change. One area in which the DLA reserves the ability to drive efficiency, and ultimately cost, is workload management. Currently, the DLA's ordering execution workload is concentrated into large, erratic monthly surges followed by rapid drop-offs. In order to prevent backlogs in order execution, the staffing level must be commensurate with the short-lived surges, creating an overstaffed condition for the majority of the month.

#### **B. ANALYSIS**

What can the DLA do to improve order efficiency? How can it reduce operating costs and improve customer satisfaction? To answer these questions, we utilized the following process to gain a clear understanding of the status quo, to explore and compare various workload-smoothing models, to use quantitative and qualitative measures to determine the most efficient model, and to make recommendations.

##### **1. Collect DLA Demand Data and Workforce Information**

We acquired five years' accumulated demand and order execution data from the DLA. Additionally, we collected literature and metrics on the agency's workforce management and workload outlay.

##### **2. Evaluate Current Order Execution Strategy**

Through a review of literature and metric data provided by the DLA, we obtained a clear understanding of the daily, monthly, and annual workload for its aviation buying staff. We then graphed and charted this data for analysis and later comparison against alternatives.

### **3. Generate Hypothetical Workload Models**

We created alternative ordering frequency models using 5-day, 10-day, and 30-day moving averages to smooth workload over the monthly period. Utilizing Little's Law, which relates inventory levels to throughput rate and cycle time, we determined optimal staffing levels required to manage each average workload level. This simulated various alternative frequencies for order execution (Little, 1961).

### **4. Quantify Impacts of Each Model and Compare to Status Quo**

We charted and graphed both status quo and hypothetical models against each other for clear visibility of both qualitative and quantitative comparisons. We predicted quantitative and qualitative impacts of each model. We generated mathematical tables to highlight differences in personnel needs (fixed costs) due to daily, weekly, monthly, and annual workload differences between various models.

### **5. Make Recommendation to the DLA Based on Efficient Staffing**

Ultimately, the DLA wishes to realize cost savings by creating efficiencies in its order execution/management plan. Currently, major workload surges require staffing in excess of normal workload needs in order to maintain appropriate service levels. With this research, we make a chief recommendation for a new order execution model based on the most efficient use of personnel, which will likely result in lower staffing requirements, and consequently, lower fixed costs. Our recommendation is based on the most efficient use of manpower and encompasses sound quantitative and qualitative analysis.

## **IV. ANALYSIS**

### **A. INTRODUCTION**

The primary objective of this project is to identify alternative strategies to minimize the DLA's forecasting and execution workload variability over the course of a given month, while still meeting customer requirements. We used the process detailed in the Chapter III to analyze and understand current operations at the DLA, then explored alternative strategies to make a recommendation on the most efficient strategy.

### **B. CURRENT STRATEGY—THE “STATUS QUO”**

The initial steps of our analysis were to collect the DLA's demand data and workforce information and then evaluate its current order execution strategy. The staffing model currently employed at the DLA seems to support a large amount of labor hours for tasks outside of order forecasting and execution. As pointed out earlier, a recent manpower study conducted internally at the DLA showed that nearly 40% of all clock time from the buying staff is dedicated to activities other than procuring spare parts for the DLA's customers. Activities such as training, daily meetings, breaks, and paid time off (PTO) represent a significant amount of lost value-added work time. Assuming these activities are required by the larger organization, in order to combat the value-added time loss, an organization must either decrease the level of work required (output) or increase the number of staff to meet the required workload. In the DLA's case, the workload amount is not adjustable, so an increase in staffing seems to be the answer.

The DLA employs 708 personnel whose specific job is to procure aviation line items for DLA's customers. Each employee works an average of 2,088 hours per year, totaling 1,478,304 works hours per year for the aviation purchasing staff. The DLA estimates that approximately 61% (1,280 hours per year) of each employee's work time is attributable specifically to purchasing activities, whereas annual training, meetings, PTO, and other administrative activities account for the remaining 39% (808 hours per year). Approximately 902,000 hours are spent on purchasing activities per year. Figure 4 shows the relationship of purchasing and other activities to total time.

## Annual Work Breakdown for DLA Aviation Buying Staff




Figure 4. Purchasing and Other Activity vs. Total Time

By examining monthly aviation consumable parts procurement request (PR) generation and processing, we can understand the workload demand on the DLA's aviation spare parts buying system. The current strategy at the DLA is to execute the vast majority purchase orders at the start of each month with some intermediate purchase orders placed during the remaining days of the month. PR processing at the DLA requires a large enough staff to adequately cover the monthly demand peaks. Under the current purchasing strategy, the buying staff must be able to handle a maximum daily workload of approximately 16,000 PRs, only to see that workload decrease to about 4,000 PRs a few days later, and even further to approximately 1,500 at the end of the month. The daily standard deviation is 3,081 PRs for the current business model. This huge workload variability is a product of the monthly business cycle in which forecasts are conducted and PRs are executed based on those monthly forecasts. The following chart (Figure 5) illustrates the current high variability in purchasing activity across a month's time. Large, short-lived PR generation activity is shown at the beginning and middle with relatively low and steady demand for the remainder of the month. To maintain customer

satisfaction under this strategy, one can see that staffing must be commensurate with the spikes to prevent backlogs, which creates the aforementioned overstaffed condition for most of the month.

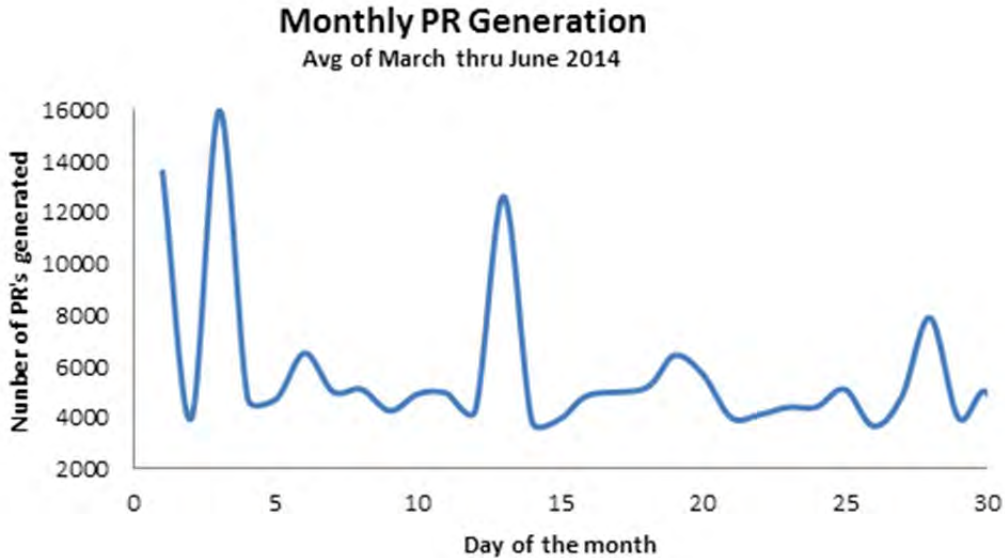


Figure 5. March through June PR Generation

The peak demand for daily PR execution per buyer is approximately 23 per day. At nearly 16,000 PRs per day, each of the 708 buyers would be required to execute approximately 23 per day. The low demand of only 1,500 PRs per day near the end of the month requires each of the buyers to execute just over two orders per day. Since the DLA does not maintain a backlog of customer order execution, we assume that its staff of 708 buyers is capable of handling this workload strategy. We assume that the average buyer is capable of executing a maximum of 23 PRs per day in a steady-state demand system.

### C. EXPLORING ALTERNATIVE STRATEGIES

The next step of our analysis required us to generate hypothetical workload alternatives for later comparison against the current strategy. Using moving averages, we simulated spreading the monthly workload out, decreasing the surges, and increasing the lulls into a smooth, consistent workload. Moving averages are a simple and effective way to smooth data. By taking the mean of a preceding set of data (e.g., the previous 5, 10, 15,

30 days) instead of the individual data points, the result is a much smoother trend. As the number of data points increases, so does the smoothness of the trend. This new trend can be followed as a means of forecasting succeeding data points or, in the DLA's case, as a new workload model for succeeding months of daily PR execution. Using moving averages, we can essentially smooth out the peaks and valleys of the DLA's current system into a more consistent workload outlay. Using moving averages maintains the original sum of the data, so the total workload for the period remains unchanged. Using a sample period (March through June 2014) of actual, and typical, workload demand data from the DLA, we could apply the moving averages and determine smoother future workload strategies that could be applied to future monthly periods.

### 1. Five-Day Moving Average

Utilizing the preceding five days of PR demand data, we added these data points and divided the sum by the number of data points (five), and we did this for each of the 30 days of the month. The result is that the workload of each day of the month is now the average of the preceding 5-day period. Figure 6 illustrates the 5-day moving average compared to the original data.




Figure 6. Current vs. 5-Day Moving Average

Figure 6 reveals that the 5-day moving average has marked reductions in variability near the beginning, middle, and end of the month, and slight increases in the ultra-low demand areas that typically follow the spikes. The overall result is a flatter line, which implies a more predictable and manageable workload outlay. At a minimum, the workload outlay utilizing a 5-day moving average would be much more manageable from a workforce management perspective.

## 2. Ten-Day Moving Average

Just as we did in the 5-day moving average, we summed the preceding 10 days of PR demand data and divided the sum by the number of data points (10); again, we did this for each of the 30 days of the month. The result was a similar but smoother daily workload than in the 5-day moving average example. The smoother trend is a result of a reduced sensitivity to trends of the original data set, which we explain later. Figure 7 compares the 10-day moving average smoothing trend to the original data.




Figure 7. Current vs. 10-Day Moving Average

Figure 7 reveals that the 10-day moving average further reduces variability (especially near the beginning, middle, and end of the month) and increases the ultra-low demand periods that typically follow the spikes. The overall result is a much flatter line—a highly predictable and easily manageable workload outlay. More so than the 5-day moving average, this alternative would facilitate an ideal personnel management strategy and may lead to easier and more consistent customer order execution.

### 3. Thirty-Day Moving Average

The 30-day moving average is the simplest of the models: a sum of the previous month’s data divided by number of days in the month. We used 30. The result is also the simplest of the alternative strategies—a flat line. Figure 8 compares the 30-day moving average smoothing to the original data.




Figure 8. Current vs. 30-Day Moving Average

### D. STATUS QUO VS. ALTERNATIVES

One might quickly deduce, with little understanding or knowledge of workload management or moving averages, that applying the 30-day moving average seems the

obvious alternative. However, there is more to it than that. As new data becomes available, it is utilized in determining the average, and the old data is dropped, thus “moving” the average along with current demand data. Based on this, it becomes clear that by increasing the length of moving average sample (i.e., 5 vs. 10 vs. 30), the resulting output becomes less sensitive to the trends of the referenced data. Figure 9 illustrates this fact by comparing the 5-day, 10-day, and 30-day moving averages to one another.

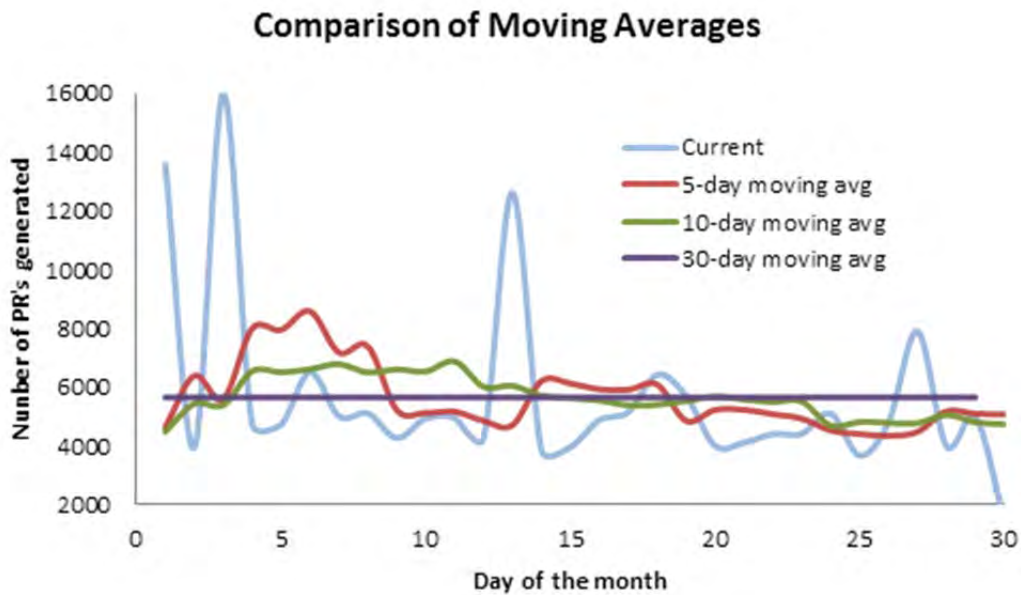


Figure 9. Moving Averages Comparison

Table 1 discloses the input data for Figures 5 through 9. The first column represents each day of a month (we use 30 days as a standard). The second column holds the original data provided to us by the DLA. It represents the average PR generation per day of the month for the months of March through June 2014. The subsequent columns are data generated in Microsoft Excel using moving averages based on the first column data. The bottom five rows are metric data (as labeled) from the above columnar data. Further explanation of each metric is provided in breakouts of Table 1.

Table 1. Input Data for Moving Average Graphs

Day	Current	5-day MA	10-day MA	30-day MA
1	13596	4649	4501	5662
2	3950	6404	5448	5662
3	15995	5611	5402	5662
4	4741	8010	6559	5662
5	4734	7960	6521	5662
6	6530	8603	6626	5662
7	5021	7190	6797	5662
8	5120	7404	6508	5662
9	4277	5229	6620	5662
10	4948	5136	6548	5662
11	4966	5179	6891	5662
12	4285	4866	6028	5662
13	12634	4719	6062	5662
14	3852	6222	5726	5662
15	3975	6137	5637	5662
16	4894	5942	5561	5662
17	5165	5928	5397	5662
18	6440	6104	5412	5662
19	5718	4865	5544	5662
20	4013	5238	5688	5662
21	4126	5246	5594	5662
22	4415	5092	5510	5662
23	4426	4942	5523	5662
24	5121	4540	4702	5662
25	3681	4420	4829	5662
26	4822	4354	4800	5662
27	7912	4493	4793	5662
28	4003	5192	5067	5662
29	4990	5108	4824	5662
30	1518	5082	4751	5662
<b>Sum</b>	<b>169868</b>	<b>169868</b>	<b>169868</b>	<b>169868</b>
<b>Mean</b>	<b>5662.27</b>	<b>5662.27</b>	<b>5662.27</b>	<b>5662.27</b>
<b>Std Dev</b>	<b>3081.23</b>	<b>1145.74</b>	<b>710.77</b>	<b>0.00</b>
<b>Maximum</b>	<b>15995.00</b>	<b>8603.20</b>	<b>6891.20</b>	<b>5662.27</b>
<b>Minimum</b>	<b>1518.00</b>	<b>4353.80</b>	<b>4501.40</b>	<b>5662.27</b>
<b>Range</b>	<b>14477.00</b>	<b>4249.40</b>	<b>2389.80</b>	<b>0.00</b>
<b>Max daily/buyer</b>	<b>22.59</b>	<b>12.15</b>	<b>9.73</b>	<b>8.00</b>
<b>Min daily/buyer</b>	<b>2.14</b>	<b>6.15</b>	<b>6.36</b>	<b>8.00</b>

Table 2 focuses on the sum, mean, and standard deviation figures, which enables us to highlight some details and advantages of moving averages. First, the sum and the mean remain constant, implying that total monthly workload (requests for the DLA’s services) does not change. However, standard deviation figures show significant range among the various alternative strategies, over 3,000 PR difference in standard deviation.

Table 2. Sum, Mean, and Standard Deviation

	Current	5-day MA	10-day MA	30-day MA
<b>Sum</b>	<b>169868</b>	<b>169868</b>	<b>169868</b>	<b>169868</b>
<b>Mean</b>	<b>5662.27</b>	<b>5662.27</b>	<b>5662.27</b>	<b>5662.27</b>
<b>Std Dev</b>	<b>3081.23</b>	<b>1145.74</b>	<b>710.77</b>	<b>0.00</b>

In Table 3, we break out the maximum, minimum, and range of the data sets from Table 1. Here, we can see the inverse relationship between the maximum and minimum figures as we move across the alternative strategies. The maximum figures of the data sets go down as we move from current strategy across the 5-day, 10-day, and 30-day moving average alternatives, while the minimum figures rise across the same. This simple revelation of the data shows how applying a moving average alternative as an ordering execution strategy would smooth out the peaks and valleys, or reduce the highs and increase the lows.

Table 3. Maximum, Minimum, and Range

	Current	5-day MA	10-day MA	30-day MA
<b>Maximum</b>	<b>15995.00</b>	<b>8603.20</b>	<b>6891.20</b>	<b>5662.27</b>
<b>Minimum</b>	<b>1518.00</b>	<b>4353.80</b>	<b>4501.40</b>	<b>5662.27</b>
<b>Range</b>	<b>14477.00</b>	<b>4249.40</b>	<b>2389.80</b>	<b>0.00</b>

Lastly, and most revealing of the advantage of moving averages, are the buyer workload figures. Breaking out these figures from the data sets in Table 1, we again see an inverse relationship between the daily maximum workload per buyer and the daily minimum workload per buyer in Table 4. As we increase the length of the moving

average data sample, the result is a significant reduction in daily maximum workload and an increase in daily minimum workload per buyer.

Table 4. Maximum and Minimum Daily Workload per Buyer

	Current	5-day MA	10-day MA	30-day MA
<b>Max daily/buyer</b>	<b>22.59</b>	<b>12.15</b>	<b>9.73</b>	<b>8.00</b>
<b>Min daily/buyer</b>	<b>2.14</b>	<b>6.15</b>	<b>6.36</b>	<b>8.00</b>

Furthermore, we graphically present the changes in daily maximum and minimum workload per buyer in Figures 10 and 11. To appreciate the stark changes in workload variability, we maintain the same scale on the x-axis of both graphs (0–25). Notice that as we move from left to right across the alternative strategies, the difference between maximum and minimum daily workload is reduced. Thus, there is less workload variability, a major objective of this project. In fact, at the 30-day moving average strategy, there is no daily variability; both maximum and minimum workload per buyer is the same. In both Figures 10 and 11, it is clear that reducing workload variability through any of these three alternative PR execution strategies would have a significant effect on staffing requirements.

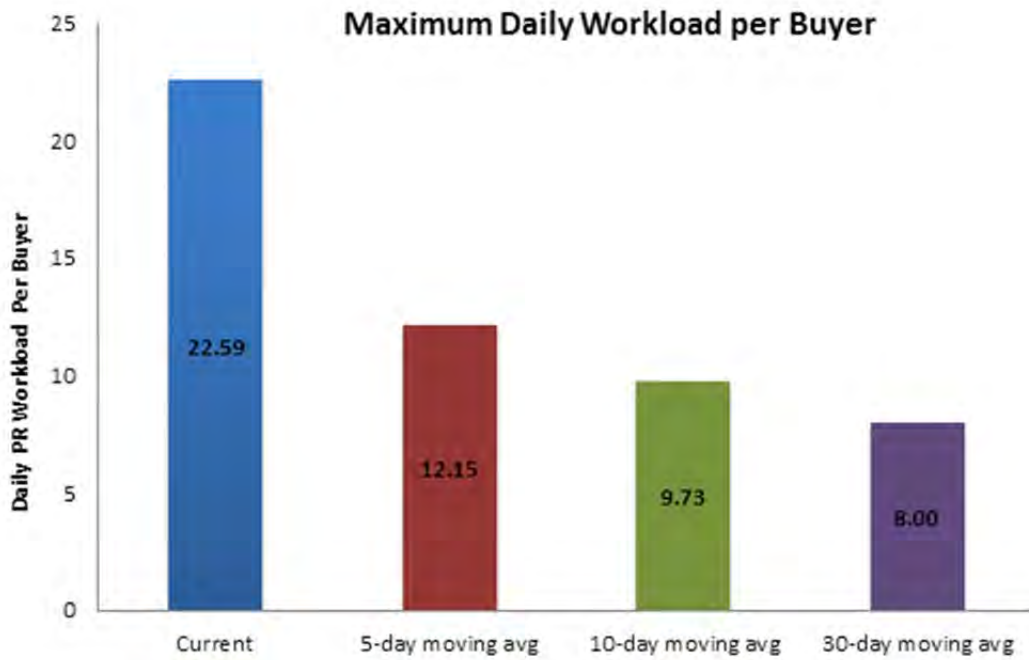


Figure 10. Maximum Daily Workload per Buyer

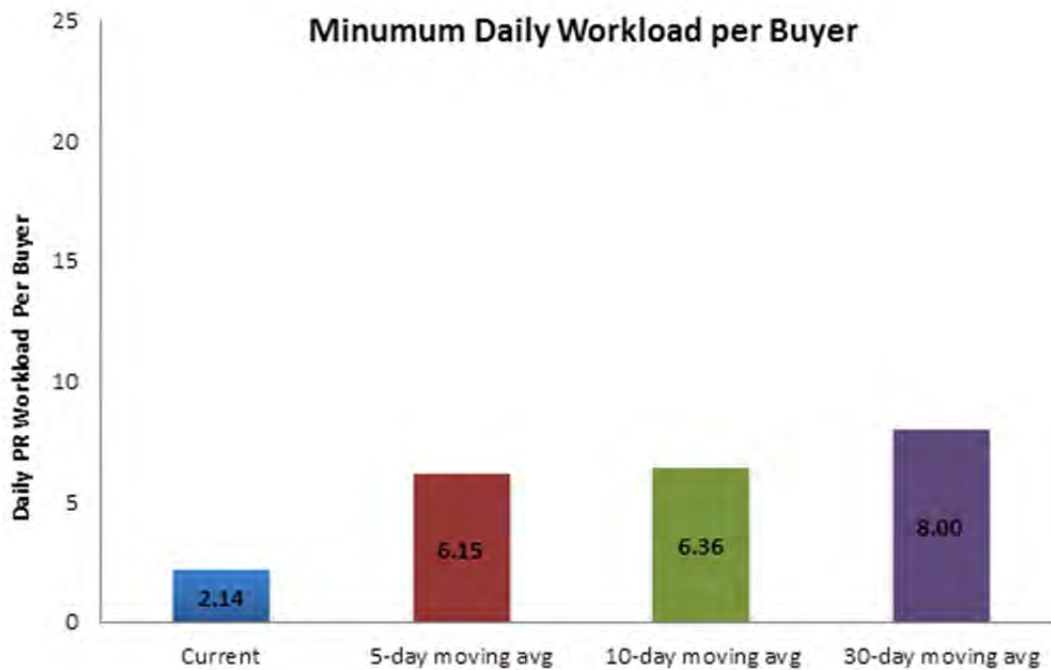


Figure 11. Minimum Daily Workload per Buyer

## E. ALTERNATIVE STAFFING REQUIREMENTS

With three alternative strategies for PR execution at the DLA identified and illustrated, we can apply Little's Law to the smoothed results of the moving averages to determine ideal staffing levels for the DLA (Little, 1961). Since Little's Law always holds true on average in a steady-state demand environment, it is critical that the surges and lulls are smoothed using moving averages before applying Little's Law to our data. In Table 5, we calculated the staffing requirement for each of the alternative strategies. We divided the maximum total daily PR output (row 4) of each strategy by the fixed maximum daily output per buyer in row 3 (PRs/buyer/day) as noted in the last sentence of Paragraph B.

Table 5. Buyer Requirements

	Current	5-day MA	10-day MA	30-day MA
Buyers (min req'd)	<b>708</b>	<b>374</b>	<b>300</b>	<b>246</b>
PR's/buyer/day	23	23	23	23
Max daily PR's	15995.00	8603.20	6891.20	5662.27

In Figure 12, we graphically illustrated the dramatic reduction in staffing requirements for each of the alternative PR execution strategies. Note that moving from the current strategy to the 5-day moving average nets a 47% reduction in staffing requirements. This is almost half of the number of buyers needed to maintain order execution and, thus, customer satisfaction. Moreover, the 10-day and 30-day moving average strategies net 57% and 65% reductions in staffing requirements, respectively.



Figure 12. Minimum Buying Staff

## F. QUANTIFYING ALTERNATIVES

If we quantify these figures in terms of annual work hours, and more specifically, work hours dedicated to buying activities, we see massive potential for fixed cost savings for the DLA. At current, each employee works an average of 2,088 hours per year, totaling 1,478,304 works hours per year for the 708-member aviation purchasing staff. With the 5-day moving average strategy, the work hours per year are reduced to just 781,000 by reducing the minimum staffing requirement to 374 buyers. Looking at the 10-day and 30-day moving average alternatives, the DLA would see annual labor hours decrease even further, to 625,600 and 514,000 hours, respectively. In terms of work hours dedicated specifically to purchasing activities, the current figure is approximately 902,000 hours annually. The 5-day, 10-day, and 30-day moving average alternatives net approximately 476,400, 381,600, and 313,500 hours of buying activity, respectively. When compared to the current strategy, these alternatives garner massive savings in labor

and personnel costs. Figure 13 illustrates the annual work hours data points discussed in this paragraph.

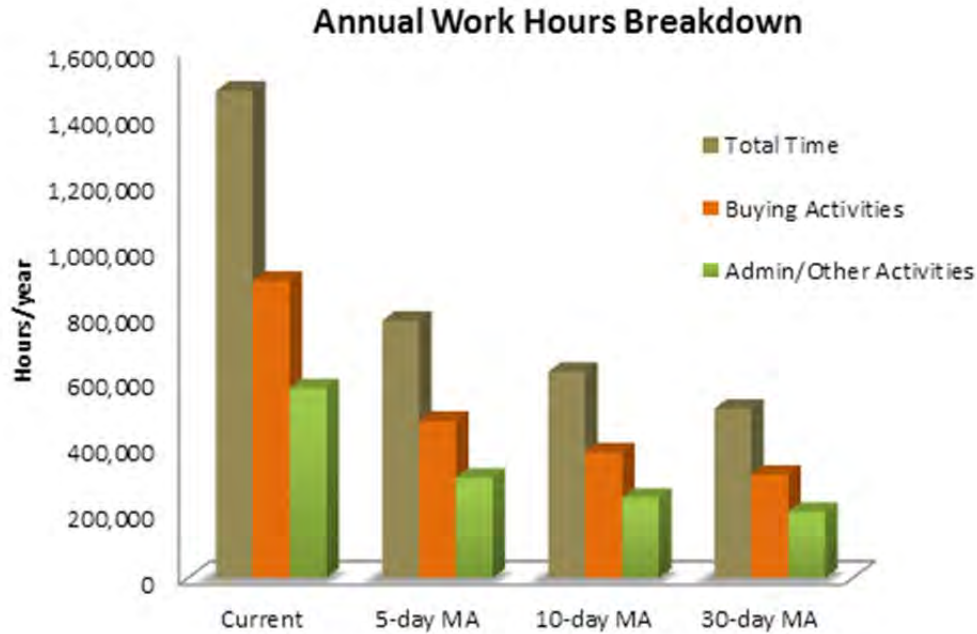


Figure 13. Annual Work Hours Breakdown

We take this one step further and monetize these alternative purchasing strategies in terms of annual labor cost savings. We calculated an average of the annual pay of the DLA’s aviation buying staff to be \$66,927 (OPM, 2014). This pay figure represents an average annual labor cost per person in the DLA’s aviation buying staff. We then multiplied this figure by the number of aviation buyers (708) to determine the current total annual labor cost to be \$47.38M. When this annual per person labor cost is applied to the alternative purchasing strategies, significant labor savings are realized. Figure 14 below illustrates the annual labor cost associated with the current system and the alternative purchasing strategies and Table 6 details the associated data. Savings are obvious and significant; net is a minimum of a 90% reduction in labor cost with the 5-day moving average and nearly a 200% reduction with the 30-day moving average strategy.



Figure 14. Annual Labor Costs

Table 6. Annual Labor Costs and Savings Data

	Annual Labor \$	Annual Savings \$	Annual Saving %
Current	\$ 47.38	N/A	N/A
5-day Moving Average	\$ 25.03	\$ 22.35	89%
10-day Moving Average	\$ 20.08	\$ 27.30	136%
30-day Moving Average	\$ 16.46	\$ 30.92	188%

Table \$ values in millions

**G. IDENTIFYING IMPACTS OF ALTERNATIVES**

The impacts of reduced staffing as a result of our proposed alternative PR execution strategies are certain and clear: reduced fixed costs, less administrative costs and lost value-added time. Additionally, the DLA will enjoy a smoother, more consistent management of work effort due to daily workload variability being greatly reduced or virtually eliminated. We anticipate the DLA can easily derive specific cost savings information by applying its labor and personnel cost metrics to our data for each of the proposed alternatives. This will produce clear fiscal data for the DLA to support change.

## **H. SUMMARY**

The primary objective of this project was to identify alternative strategies to minimize the DLA's forecasting and execution workload variability. In this chapter, we began by understanding and evaluating the current system in place at the DLA—the status quo. Through this, we learned that the DLA's current system allows for a significant amount of non-value added clock time and the staffing model is centered around a single monthly workload surge. With a clear understanding of how the DLA does business at current, we set to researching and generating alternative workload management strategies. We developed three alternative strategies using the concept of simple moving averages and compared these to the status quo. We learned that by applying these strategies to the DLA's current monthly customer demand, we could dramatically reduce the staffing level required to handle the workload, and as such, reduce the DLA's fixed operating costs and ease workforce management.

## **V. RECOMMENDATIONS AND CONCLUSIONS**

### **A. RECOMMENDATIONS**

We showed in our analysis that workload smoothing holds tremendous positive benefits in terms of efficiency and staffing utilization. Having never worked at the DLA, we do not pretend to be experts at its business. That being said, there is certainly room for efficiency improvements within any organization. We believe that order execution could and should be spread out over the month in a more distributed fashion than currently exists at the DLA. Rather than being so bold as to suggest sweeping and immediate changes to a business model that has been used effectively at the DLA for a long time, our recommendations for implementation are cautious and incremental.

Organizational change is never easy, and getting buy-in from the leadership and down is key. We suggest testing out the alternative strategies we presented regarding workload smoothing on a small segment of the business at the DLA. The selection of this pilot group will be the key to future success. Ideally, the pilot group will be (1) receptive to new ideas and willing to change its business model, (2) small enough in size so as to be manageable and tightly controlled, (3) established enough to form a meaningful baseline for what will be somewhat of an experiment, and (4) responsible for items with less critical delivery dates so that the failure risk is minimized.

Once the pilot group is selected, we recommend management commit to a full year of implementing an order execution strategy outlined in Chapter IV. This extended time horizon permits the pilot group to become accustomed to the new way of conducting order execution and allows time for the new process to settle out and become optimized within the workgroup so that it can be fairly compared with the status quo. During this trial year, management tracking of daily staff workloads will be important as this data will be used to compare the old and new models to quantify efficiency gains captured by the new business model. As we showed in our analysis, the more evenly distributed the workload, the greater the potential for efficiency gains. The goal is to reduce workload variability as much as possible. The workload spikes generated by forecast updates can be smoothed by forcibly reducing the workload to approximately that which is

represented by the 30-day historical moving average. This will cause orders to be delayed by as much as two weeks so it is not a suitable methodology for the most urgent needs. Truncating workload will have some measurable impact on customer service levels that needs to be quantified in terms of forecast degradation or number of stock-out occurrences during the experimental phase as compared with the baseline, and this is why we recommend a small-scale initial implementation.

During the course of the one-year trial, management can track progress on a monthly basis of the 30-day moving average workload model to determine if the positive effects outweigh the costs incurred (costs coming in the form of increased chance of stock out). Specific metrics might include order timeliness, forecasting quality, workload, and workload variability during the course of each month. This data can be compared each month and, more meaningfully, at the end of the year-long trial period.

## **B. CONCLUSIONS**

Performing the voluminous PR processing at the DLA requires a large enough staff to adequately cover the overall monthly demand. The current staffing is large enough that it covers the demand peaks experienced at the beginning and middle of each month. While this arrangement satisfies the required demand peaks, it potentially leaves some workforce elements underutilized during periods of lower demand. We believe our analysis has shown that the order execution workload variability experienced by the DLA is inefficient and possibly avoidable. If this workload variability can be reduced or eliminated there is an opportunity for significant savings in labor costs. We understand that there exists a monthly business cycle with fluctuating workload. By smoothing the workload over the course of a month, it is possible to reduce average monthly staffing requirements significantly. We have shown the positive correlation between workload smoothing and reduced labor requirements. The challenge ahead for the DLA lies in changing the culture of the organization so that the workload smoothing is accomplished while maintaining an acceptable customer service level.

## LIST OF REFERENCES

- Altay, N., & Litteral, L. A. (2011). *Service parts management demand forecasting and inventory control*. New York, NY: Springer.
- Cecere, L. M. (2013). *Bricks matter: The role of supply chains in building market-driven differentiation*. Somerset, NJ: Wiley & Sons.
- Crum, C., & Palmatier, G. E. (2003). *Demand management best practices: Process, principles, and collaboration*. Plantation, FL: J. Ross.
- Defense Logistics Agency (DLA). (2013a). *Demand planning overview—The basics* [Job aid]. Fort Belvoir, VA: DLA.
- Defense Logistics Agency (DLA). (2013b, October). DLA at a glance. Retrieved from <http://www.dla.mil/Pages/ataglance.aspx>
- Defense Logistics Agency (DLA). (2014). *Demand classification job aid*. Fort Belvoir, VA: DLA.
- Edwards, J. E. (2011). *DOD's 2010 Comprehensive Inventory Management Improvement Plan addressed statutory requirements, but faces implementation challenges* (GAO-11-240R). Washington, DC: Government Accountability Office.
- Johnson, N. B. (2013, August). *DLA director calls for \$13B in budget cuts*. Retrieved from <http://www.federaltimes.com/article/20130825/AGENCY01/308250002/DLA-director-calls-13B-budget-cuts>
- Laha, D., & Mandal, P. (2008). *Handbook of computational intelligence in manufacturing and production management*. Hershey, PA: IGI Global.
- Little, J. D. C. (1961). A proof for the queuing formula:  $L = \lambda W$ . *Operations Research*, 9(3): 383–387.
- Pegels, C. C. (2005). *Proven solutions for improving supply chain performance*. Greenwich, CT: Information Age.
- U.S. Office of Personnel Management (OPM). (2014). Pay & leave: Salary & wages: Salary Table 2014-RCH. Retrieved from <http://www.opm.gov/policy-data-oversight/pay-leave/salaries-wages/salary-tables/14Tables/html/RCH.aspx>
- Wallace, T. (2006, Spring). Forecasting and sales and operations planning: Synergy in action. *The Journal of Business Forecasting*, 25(1):75–85.

THIS PAGE INTENTIONALLY LEFT BLANK

## **INITIAL DISTRIBUTION LIST**

1. Defense Technical Information Center  
Ft. Belvoir, Virginia
2. Dudley Knox Library  
Naval Postgraduate School  
Monterey, California