

Computing and Visualizing Decision Impact in Enterprise Resource Planning Systems

Final Report for Microsoft Contract #7440469

Prepared by:

Stephen F. Smith (PI), Charles Buhman, Willem Jan van Hoeve,
Pradeep Varakantham and Yangfang Zhou

Tepper School of Business and The Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh PA 15213

With User Interface Design Contributions from:

Morten Holm-Petersen
Microsoft Corporation

April 2009

1 Introduction

In this final report, we summarize research carried out by Carnegie Mellon University in collaboration with Microsoft Corporation under contract #7440469, titled *Computing and Visualizing Decision Impact in Enterprise Resource Planning (ERP) Systems*, over the period from March 2008 to March 2009. The broad focus of this research has been to enhance role-tailored decision-support within Enterprise Resource Planning (ERP) systems such as Microsoft Dynamics AX. Although current ERP systems contain much of the information necessary to support coordinated role-tailored decision-making, they continue to be used to support decision-making in compartmentalized ways and provide very little in the way of advanced analytics for conveying the larger impact of potential decisions. For example, a product manager accepts new orders and negotiates order deadlines based on prior order data accessible to this role and does not consider data relating to current production commitments and constraints. A production manager takes inventory management decisions based on current production and supplier delivery data but does not consider planned marketing campaigns. Such compartmentalized decision-making inevitably leads to suboptimal behaviors.

To address this problem, our research has investigated the incorporation and use of modern optimization tools. Our claim is that optimization models that capture the key inter-dependencies that link different decision-making roles and business processes can be used (1) to provide decision-makers with a global picture of the consequences of important business decisions on different components of the organization, and 2) to allow decision-makers to rapidly assess the overall impact of different decision options; hence enabling decisions that ultimately lead to more optimized organizational performance. We further claim that such optimization models can be driven by data that is derived directly from the information in the underlying ERP system, and hence can be layered directly on top of existing ERP capabilities. To integrate optimization models gracefully for role-based decision support, we propose the development of an explicit business process representation layer. By linking optimization models directly to decision points in role-based business processes [MS 2007], we believe such models can be exploited in a natural way that does not require any special user expertise in advanced analytics or optimization.

Our broader vision of the research is a graphical business process “control center” that is layered over the basic role-based interface and functionality of a contemporary ERP system such as Microsoft Dynamics AX. Imagine a computer screen that depicts the critical business processes being managed within a given role, together with linkages to other inter-dependent business processes. Imagine further that these processes are instrumented with sensors (analogous to push pins on a network diagram). Sensors provide a continuous, graphical basis for tracking the performance and state of uncertain aspects of business processes, and for detecting the need for action or change. Zooming on a sensor exposes relevant ERP system data, abstracted in the form of key business metrics. Underlying analysis and optimization models are applied when the graphical

business process map is interrogated to show the current upstream/downstream impact of local decisions and processes on business processes owned by other roles. These analysis and optimization models are also accessible within the control center for use in what-if, decision-analysis mode, to convey the financial tradeoffs of different business decision options and suggest the best course of action.

Toward this vision, we have focused on developing and applying an optimization model to improve decision-making within two core interacting business processes: demand forecasting (carried out by the sales department in conjunction with marketing) and supply planning (the responsibility of the production department). Our initial research effort has produced the following results:

- *Supply planning optimization model* - Using data obtained from Hoshino USA (a NAV ERP system user) as a test environment, we developed a prototype model for solving the basic supply planning problem of deciding how much to order from suppliers when to maximize service level while minimizing inventory costs. To address with the uncertainty inherent in both demand forecasts and supply planning constraints, the model implements a stochastic optimization approach: a core constraint optimization procedure (implemented as a mixed-integer program in CPLEX) is coupled with a “scenario” generator and the results of multiple runs are combined to produce the *most likely* outcome.
- *Construction of demand probability distributions from ERP demand data* – The generation of scenarios relies on the availability of probability distributions of customer demand. To infer demand distributions from historical customer data resident in the base ERP system, we defined and prototyped an approach that couples clustering and kernel density estimation techniques from the field of machine learning. Although developed for purposes of computing demand distributions and demand forecasts, the approach is equally relevant to constructing models of the uncertainty associated with various planning constraints (e.g., supplier lead time distributions).
- *Model use scenarios and user interface design* - In collaboration with Morten Holm-Peterson of Microsoft Corp., we developed a set of model use scenarios based on Microsoft’s role/persona framework, and then used these scenarios to produce a visionary Dynamics GUI design.

The remainder of this report is organized as follows. In Section 2, we describe the approach taken to development of the optimization model, summarize the demand and supply planning tools that have been developed, and distinguish their advantages with respect to contemporary ERP planning capabilities. In Section 3, we present a set of reference user tasks for applying the prototype tools to support demand and supply planning processes. In Section 4, we then take a subset of these reference tasks and storyboard a graphical user interface design for a subset of these reference tasks. In Section 5 we summarize the unique opportunities offered by modern optimization techniques in extending the functionality of contemporary ERP systems, and Finally in Section 6 we discuss next steps. In Appendix A, a formal specification of implemented CPLEX supply planning model is given, and Appendix B contains the set of UI screen images developed by Morten Holm-Petersen that are referenced in Section 4.

2. Scenario-based Supply Planning

As indicated above, we take a *scenario-based* approach to optimization. There are two basic reasons:

- There is significant uncertainty in manufacturing operations and effective decision support must take this into account. Scenario-based reasoning is a form of stochastic optimization that utilizes explicit models of uncertainty (in the form of probability distributions) to produce outcomes that more accurately reflect actual operating conditions
- Modern constraint programming techniques provide powerful and flexible tools for solving complex combinatorial optimization problems. Scenario-based reasoning provides a framework for coupling such deterministic solving techniques with models of domain uncertainty.

For these sorts of reasons, the use of scenario-based reasoning techniques has been gaining increasing momentum within several branches of the optimization research community in recent years (e.g., Amram and Kulatilaka 1999, Van Hentenryck and Bent 2006, Savage et. al 2006, Tarim et. al 2006).

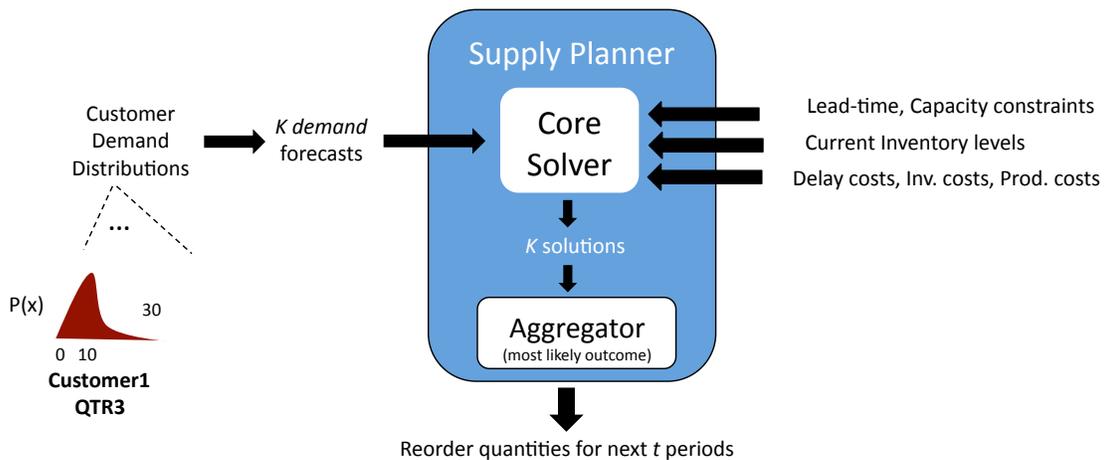


Figure 1: Scenario-Based Supply Planning Model

Our overall supply planning optimization framework is graphically depicted in Figure 1. At its core is a deterministic solver built using constraint reasoning techniques. Given a demand forecast, it solves the basic problem of how much to order from suppliers when to minimize costs (and hence maximize profits). This core solver is embedded in a larger stochastic optimization loop, where k demand forecasts are constructed from underlying customer demand distributions, the k “scenarios” are solved by the core solver, and then the results are combined (via an aggregation function) to produce the most likely outcome.

2.1 Modeling Demand and Generating Demand Forecasts

Application of this framework to the supply planning problem faced by any given organization immediately raises the question of how to obtain appropriate models of customer demand. A basic assumption made by the framework is that demand can be modeled as a probability distribution, and in the simplest case, forecasts are built by just sampling these distributions. Demand distributions can be obtained in a number of ways. Since our focus is on integrating optimization capabilities within an ERP environment, our research has explored mechanisms for constructing distributions from historical data resident in the ERP system. However, distributions can also be generated or adjusted by users based on their external knowledge. Distributions can also be manipulated (e.g., skewed) to construct more “risk-taking” or “risk-averse” forecasts for evaluation.¹

Figure 2 graphically depicts the approach we have taken to building customer demand distributions from base ERP sales data. The key issue is how to extrapolate from sparse past data. Consider the following example. A manufacturer’s customer demand varies considerably from period to period and hence it makes sense to build models of demand at this level of granularity. Suppose that a period is one month. Then a customer that the firm has sold to for the past 5 years will have at most 5 data points from which to build a demand distribution for any given period (i.e., one data point for each year).

To overcome this problem, we combine two ideas from the field of machine learning:

- *Demand Clustering* – One straightforward way to generate more data points is to combine individual customers into customer classes, based on commonality along some set of criteria. In analyzing ERP data provided by Hoshino USA in support of this project, we have found it productive to cluster customer sales data by Sales Order Quantity. This criterion has the additional benefit of clustering smaller (and in Hoshino’s case) less reliable customers into equivalent classes, which allow us to produce distributions that smooth demand and further increase forecasting accuracy. However, one could imagine a range of additional clustering criteria that could provide a viable basis for generating alternative demand forecasts (e.g., OrderCount12Months, AverageProfitMarginPercentage, CustomerRegion, CustomerGroup). The kernel computational model we have employed as a demand clustering tool is a basic *k-nearest neighbors* clustering algorithm. However, there are a broad range of more sophisticated clustering algorithms that might be adopted for this purpose.
- *Kernel Density Estimation* – Once clusters have been determined, the order quantity data associated with all customers in any given cluster is used as a set of sample points from which to generate a probability distribution. To this end, a basic kernel density estimation (KDE) procedure is used. Basically, KDE works

¹ In fact, it is not strictly necessary to drive the optimization model from demand distributions. In some circumstances, users might want to build specific (non-stochastic) scenarios that reflect different possible futures.

by considering each sample value as the center of a uniform distribution, overlaying these distributions, and then summing the probability mass at each value along the x-axis to build a composite distribution. One particular advantage of KDE that argues for its use in this context is that it tends to be quite effective with relatively small numbers of data points.

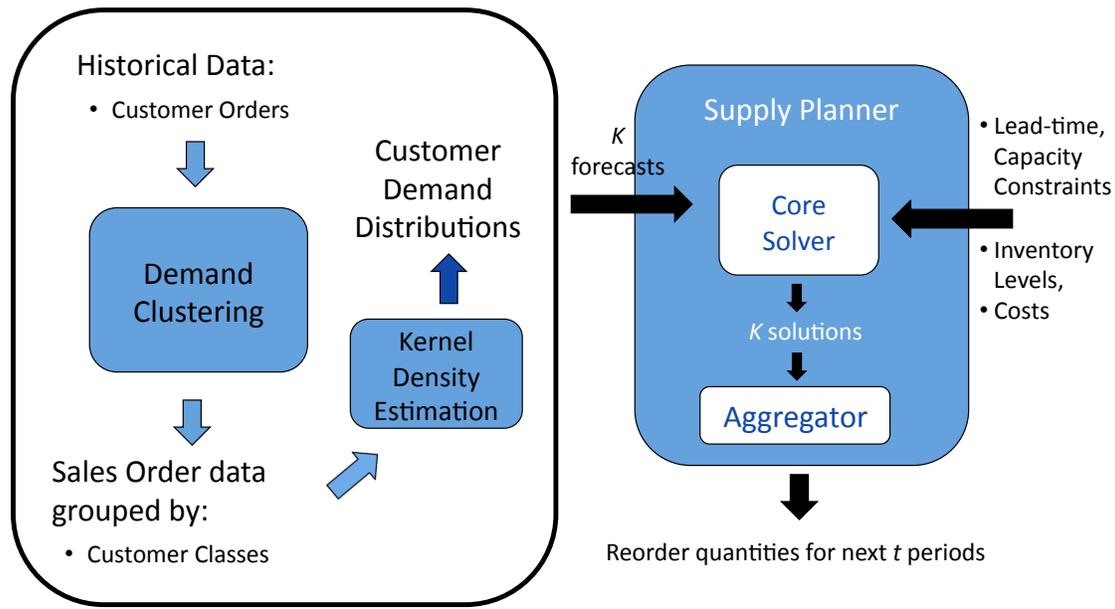


Figure 2: Constructing Forecasts from Historical Data

2.2 Core Supply Planning Model

The supply planning model that has been developed as the core solver is an aggregate, period-based model for determining what quantities to order from suppliers when to minimize costs (maximize profits). A period can be set to any appropriate value (e.g., a day, a week, a month) and establishes the overall granularity of the aggregate model. It is expected that the value a period takes on will be different in different application contexts. In the case of Hoshino USA, for example, reordering decisions were made once a month and hence a period length of a month is logical.

The supply planning model takes four basic inputs:

- A *demand forecast* for the next n periods for all products and customers. In some settings it may make more sense to consider demand at the level of product groups and customer groups, or for some combination of individual products/customers and product/customer groups. (We will use the terms customer and product interchangeably with customer group and product group in what follows.) A customer demand forecast for any particular product is assumed to be a set of n discrete order quantity values (one for each period). Recall from above that an input forecast corresponds to a “scenario” in our framework and is typically constructed by sampling an underlying probability distribution.

- *Inventory levels* - An indication of current inventory for each product, both actual inventory for the current period and expected inventory (i.e., already ordered and in-transit) for the next $n-1$ future time periods.
- *Lead time constraints* on supplier orders and production of final products. The model currently assumes a single lead-time constraint that encapsulates entire process of producing and shipping supplier parts, assembling and or manufacturing the customer product and transport time to the customer. The lead-time constraint indicates how many periods it will take to acquire component parts, assemble/manufacture final products and deliver to customer. There is no conceptual reason why we couldn't have separate lead-time constraints for supplier orders, in house assembly/manufacturing and customer delivery if the latter provides a better story. In the case of Hoshino data we have used as a reference case, it is supplier lead-time that determines when customers get orders satisfied (there is a small in-house kitting operation but this time is negligible in comparison).
- *Capacity constraints* - Each supplier has a maximum production capacity per period, which limits the amount of product that can be ordered per period. The model also includes transport container capacity constraints, to account for the fact that products are transported in batch quantities.
- *Costs* - A final set of inputs specify relevant costs include inventory costs (per product) and delivery delay costs (per customer and product).

The model generates two principal outputs:

- What product quantities to order from suppliers in each of the next n periods. Only the reorder quantities generated for the next period are considered actionable data.
- An expected allocation of current and future inventory over the next n periods to various customers to best meet demand.

In allocating supply to demand over the next n periods, the optimization objective of the model in the basic (unconstrained) case is to maximize profit, or equivalently to minimize inventory and customer delay costs. Minimization of customer delay costs corresponds directly to maximizing customer service level. The model can also be constrained to achieve a minimum service level (either globally or with respect to particular customer classes), in which case the optimization additionally respects designated service level constraints.²

² Note that the supply planning model is maximizing customer service level in an aggregate sense, determining how to best satisfy demand on a per period basis. To actually assign inventory to known orders in the current period and evaluate service level performance, we would propose use of a separate allocation procedure. In this short-term setting a deterministic procedure that reasons with expected values for various duration constraints is likely more appropriate. We would still advocate a constraint-based search procedure, given their general effectiveness in combinatorial scheduling and allocation

Another simple parameter in the basic supply planning model that can be manipulated by the user is the relative emphasis to be placed on customer service versus inventory cost. Essentially, the model assumes that the overall objective is built as a weighted sum, and the user might decide to de-emphasize inventory cost (e.g., if firm is in a growth mode) or de-emphasize customer service level (e.g., if firm is strapped for cash).

To illustrate the versatility of such a supply planning model to support different decision-making roles, an extension was designed to support a marketing function. Specifically, a so-called *market optimization* tool was created by adding one more input parameter – an amount of additional capital to invest in a marketing campaign – and solving the same problem as before to determine which products and customers to focus on to maximize overall profit. The same customer preference ratings that are used to minimize customer delay costs together with the price/cost differential associated with various products now drive the determination of how to best spend the additional capital.

2.3 Implementation

A prototype implementation of the scenario-based supply planning framework was developed and demonstrated using sales and purchase order data provided by Hoshino USA. The core supply planning optimization model was implemented as a Mixed-Integer Program using ILOG's CPLEX solver. This modeling choice was based principally on the structure of the supply planning model that emerged from our formulation, which in turn was influenced by the nature of Hoshino USA's supply planning problem. Hoshino (the parent company) is a leading manufacturer of guitars and drums. Hoshino USA acts principally as the US distributor for the parent company back in Japan, and as such, its supply planning problem lies principally in minimizing inventory costs and maximizing service level. There is some amount of kitting (e.g., of amps and guitars) and repackaging that gets carried out on-site, but overall there is no in-house production activity.

If the supply planning problem to be solved had alternatively involved a significant internal production planning component, with its associated capacity and scheduling constraints, it is likely that we would have instead implemented using a constraint programming (CP) tool like ILOG solver, since for these types of combinatorial problems, CP has proved to be a very powerful and scalable technology. As future research pushes toward a more general supply planning optimization model, we expect that a CP-based approach will be a better way to go.

To test the prototype supply planning tool, Hoshino provided us with 3 years of customer sales data and supplier purchase order data for 3 SKUs that provide different demand forecasting and supply planning challenges:

- A highly seasonal guitar package
- A proprietary model of Hoshino USA's largest customer with the property that its supply lead time is greater than its customer order lead time
- A new model whose demand is highly uncertain

applications. But presumably a classical ERP production planning capability could also be applied here.

Experiments were performed to confirm that application of clustering according to sales volume and kernel density estimation to the first 2 year's data resulted in demand distributions representative of the year 3 data. A similar experiment was performed to gauge the potential of the supply planning tool. However, idiosyncratic constraints relating to Hoshino's credit practices with unreliable customers prevented us from making a meaningful assessment of how much better performance could have been if the tool was in use.

A formal specification of the basic supply planning model is given in Appendix A.

3. Scenario Tasks

In this section, we develop a sequence of linked demand forecasting and supply planning tasks to illustrate how we envision the clustering and optimization tools described in the last section integrating with and enhancing relevant business processes. In describing these tasks, we adopt the role/persona framework developed for MS Dynamics at Microsoft Corporation. Specifically we center on the roles of Kevin (the Sales Manager), Eduardo (the Production Planner) and Ellen (the Warehouse Manager).

Task 1: Produce Customer Demand Forecast

Task Goal: Determine demand forecast for next n periods for all products and all customers (or all product groups and customer groups if this level of aggregation is more appropriate)

Persona: Kevin (sales manager)

Trigger: Regularly each period (in the case of Hoshino this happens once a month)

Input Data:

- Actual customer orders for delivery in future periods
- Sales forecasts from each sales representative
- Historical customer data: Customer orders (quantity, frequency) in previous periods and same period in previous years
- Historical measure of salesperson's forecasting accuracy

Outcome: One or more customer demand forecasts (represented as discrete probability distributions - see discussion of inputs to the Supply Planning tool in Section 2), along with target customer service levels. If more than one demand forecast is generated, then each is tagged with defining assumptions.

Task Steps:

- Three days before the end of each period (month) Kevin receives emails from each sales rep with forecast numbers for future orders for the next n periods in their area. Each sales rep is asked for low, high and most likely numbers for each period, along with indications of the likelihood of each of these alternatives. These forecasts are automatically assembled into a single consolidated forecast, with the inputs of a given sales person being displayed in a way that indicates that sales rep's general forecasting accuracy. The system utilizes a given rep's historical measure of forecasting accuracy. It updates this measure each period by comparing the actual orders taken in during the previous period (by extracting appropriate data from the ERP system) to the numbers projected in the sales person's previous forecast, and bumping the sales rep's previous accuracy value up or down depending on the results of this comparison.
- The *demand clustering* tool is also automatically invoked to compute clusters using a number of pre-selected clustering criteria and shows the user the effectiveness of each in terms of smoothing the unevenness of historical demand patterns. These are displayed visually for comparison (as described in Section 1). To provide additional basis for comparison and evaluation, the demand distribution following from each clustering alternative is also generated and displayed in a correlated fashion.

- Equipped with visualizations of both the consolidated sales rep forecasts and the forecasts resulting from various historical data clustering alternatives, Kevin analyzes the similarities and differences across forecasts. He first analyzes the set of forecast distributions generated by adopting different customer data clustering criteria, looking first at the pre-selected clustering criteria and then generating another run with a slightly different configuration of clustering criteria. After subsequent comparison, Kevin decides to select one of the machine generated demand distributions as most representative of past demand patterns. This constitutes his current candidate demand forecast.
- Next, Kevin compares this forecast with the consolidated forecast produced by integrating the numbers that each sales rep. provided. Points of agreement and divergence are visualized and analyzed, using visibility of the historical accuracy of various sales persons' forecasts to drill into the details of particular line items as well as extra knowledge that Kevin might be aware of.
- Finally, manipulates the demand forecast generated from historical data to incorporate selected sales rep. forecast information and produce a final baseline demand forecast for the current period, using estimates of sales rep forecasting accuracy to decide what weight to place on sales rep versus historical forecast distribution information. He also determines a minimal target for customer service level that he would like production to achieve over the next n periods. Both the baseline demand forecast and the target customer service level are communicated to the production planning department (Eduardo and Ellen).
- (Optional) Kevin creates additional demand forecast distributions based on external knowledge that he has obtained (e.g., that a new competing product is expected to win part of the market). For each additional alternative generated, a target customer service level goal. Each additional alternative is also communicated to the production planning cell, together with an associated likelihood of occurrence.

Next task: *Determine Supplier Orders and Aggregate Supply Plan* task.

Task 2: Determine Supplier Orders and Aggregate Supply Plan

Task Goal: Given the current product inventory, the current demand forecast, and supplier lead time and capacity constraints, determine

- the quantity of each product that needs to be ordered from suppliers this period
- an allocation of products to customer demand over the next n periods

that minimizes inventory cost (maximizes profit) and maximizes customer service level (% of orders received in the period in which demand is projected).

Persona: Eduardo (production planner), Ellen (warehouse manager).

Trigger: Regularly, each period, once the task "Produce customer demand forecast" is done.

Input Data:

- One or more customer demand forecast distributions, together with target customer service level goals and likelihood of occurrence
- Current inventory, including products currently in the warehouse, and those associated with open purchase orders that are expected to arrive in future periods

(according to lead time constraints)

Outcome: A quantity for each product (or product group) that must be ordered from suppliers in the current period to meet projected demand in future periods at the final determined service level.

Task Steps:

- Eduardo and Ellen meet face-to-face and load up the current demand forecast produced and communicated by Kevin into the Supply Planning Tool. The supply Planning tool is first run to compute the optimal allocation for the current baseline weighting of inventory cost minimization and customer service level satisfaction. This solution is evaluated with respect to the target customer service level provided by Kevin (sales). If the customer service level achieved does not achieve the target, the following actions are taken.
- A range of customer service level constraints are specified (e.g., from 100% to 90% in increments of 2) and the model is re-invoked with this customer service profile to produce a graph indicating the tradeoff in inventory cost as the customer service level is decreased. Several outcomes are possible. For example, if Kevin has specified a 95% customer service level and it is observed that there is significant savings in inventory cost when moving from 95% customer service level to 94% customer service level, then Eduardo may initiate communication with Kevin in Sales to see if this is an acceptable compromise.
- For each unfilled component of the current demand forecast, an analysis is made of various possibilities for meeting this unfulfilled demand. Eduardo and Ellen first look for excess capacity for other products (or product groups) and if found in the current or future periods attempt to re-purpose that capacity to lessen expected shortfalls. Failing this, Eduardo and Ellen produce estimates of how much it would cost in capacity investments or outsourcing to meet the full extent of the current demand forecast in future periods. The Supply planning tool is run again iteratively, giving increasingly less weight to minimizing inventory cost on each run, to determine whether increasing inventory costs is a better option than outsourcing for meeting the unaccounted future demand.
- After analyzing these alternatives together with the uncertainty associated with the customer demand forecast, a decision is taken whether to prefer one of the above three options (repurpose capacity, increase inventory costs or outsource unfilled demand), or whether neither option should be preferred and the extra demand should just be left unaccounted for (at the potential of some loss of customer satisfaction).
- In circumstances where Eduardo and Ellen recognize that there is excess capacity, Kevin (in sales) is informed of the possibility for satisfying additional demand if it can be generated. Kevin applies the same model to explore these possibilities (see task 3 below) and responds with a modified demand forecast.
- After determining finalizing the supply plan to be adopted, the quantities of each product (or product group) that must be ordered in the current period is generated. This supplier order data is fed into a purchase requisition worksheet and forwarded to the Purchasing Department, where purchasers will then group demands for different products with different vendors to negotiate a discount (or consolidate with a single vendor to reach the quantity agreed in a frame

agreement that triggers a large discount).

Next Task: *Generate Supplier Purchase Orders* task.

Task 3: Analyze potential for special marketing campaign

Task Goal: Determine if it is advantageous to undertake a special marketing campaign to increase demand, and if so, determine which customers and products should be targeted.

Persona: Kevin (Sales manager)

Trigger: Indication from production planning cell that there is potential for satisfying more demand (e.g., excess capacity, excess inventory)

Input Data:

- Current demand forecast
- Current inventory, including products currently in the warehouse, and those associated with open purchase orders that are expected to arrive over time (according to lead time constraints)

Outcome: Revised demand forecast that includes additional demand expected from investment in marketing campaign

Task Steps:

- Kevin considers possibilities for further increasing profits through the introduction of a special marketing campaign. The marketing optimization tool (actually a variation of the same model used by supply planning) is consulted to identify the products/customers that will give the biggest impact for various levels of additional investment. As a first step, Kevin reviews the KPI's used to define the customer preference (priority) that is encoded in the service level delay penalty, and adjusts the weights of specific KPI factors to fit current circumstances. Among the KPIs that might be of interest in prioritized customers are: sales volume, profit margin (due to contractual terms), reliability of payment, and past history of delays (fairness). Note that this action of selecting and setting the weights of KPIs that will determine the service level delay penalties is equally relevant to the aggregate supply allocations that are made in Task 2.
- Once customer priorities (delay penalties are finalized, the model is run several times with different levels of investment. A graph is produced showing the profit cost tradeoff as increasing investment is made.
- If run in base mode, the graph produced in the task step above will reflect a risk neutral posture. Kevin may decide to adjust the risk level of the model, either to assume a more risk taking posture or to analyze more risk adverse option. In either case, the model can be run again to produce alternative results for comparison.
- In the end, Kevin determines whether a special marketing campaign is desirable and if so, what level of investment is most appropriate. A revised demand forecast is produced based on this decision and communicated back to the production planning cell.

Next Step: Back to the “*Determine Supplier Orders and Aggregate Supply Plan*” task

4.0 Storyboards for Scenario Tasks

In the following, the a subset of the scenario tasks specified in the last section are specialized to coincide with a series of GUI interface design screen images produced by Morten Holm-Petersen of Microsoft Corporation. The referenced screen images appear in Appendix B.

Task 1: Produce Customer Demand Forecast

Task Goal: Determine demand forecast for next n periods for all product and customer groups

Persona: Kevin (sales manager)

Trigger: Regularly each period - for purposes of the demo we will assume that the forecast is computed once a month

Input Data:

- Actual customer orders for delivery in future periods
- Sales forecasts from each sales representative
- Historical customer data from AX: Customer order data (quantity, frequency) in previous periods and same period in previous years
- Historical measure of salesperson's forecasting accuracy

Outcome: A demand forecast for all product and customer groups (represented as probability distributions that capture the inherent uncertainty); optionally a target service level constraint may also be specified for gold customers. For purposes of the demo we will assume that the forecast projects ahead 3 periods (months)

Task Steps:

- Three days before the end of each period (month) Kevin receives emails from each sales rep with forecast numbers for future orders for the next n periods in their area. Each sales rep is asked for low, high and most likely numbers for each period, along with indications of the likelihood of each of these alternatives. **[Slide 3]**. These forecasts are automatically assembled into a single consolidated forecast, and stored within the forecasting role center of AX for later consideration **[Slide 4]**. In producing this consolidated forecast, the system utilizes a given rep's historical measure of forecasting accuracy. It updates this measure each period by comparing the actual orders taken in during the previous period (by extracting appropriate data from the ERP system) to the numbers projected in the sales person's previous forecast, and bumping the sales rep's previous accuracy value up or down depending on the results of this comparison.
- To develop the forecast, Kevin enters the forecasting role center of AX **[Slide 4]**. Using its default settings and the sales forecast data just entered, the system computes an initial forecast. Kevin decides to first view the forecast from a product group perspective **[Slide 5]**. The display shows the current demand projection for each product group (in order of sales volume) along with the uncertainty associated with each that is implied by the underlying demand probability distributions. The demand distributions follow in part from the system's current default criteria for clustering customers into customer groups.
- Kevin clicks on the customer cluster strength indicator to examine in more detail what clustering assumptions are being made in computing this initial forecast.

- [**Slide 6**] This brings up a display that characterizes the quality of the current default customer clustering assumptions along dimensions such as order behavior consistency and consistency with respect to historical trends [**Slide 7**]. Kevin first asks the system to visually depict its current customer clusters [**Slide 8-9**]. Kevin then decides to compare this result to that obtained with the customer clustering assumptions used for prior 2008 forecasts [**Slides 10-11**]. Upon seeing a better clustering result from the standpoint of ordering volume, Kevin decides to switch to the 2008 customer clustering assumptions [**Slide 12**]. Upon clicking ok, an annotation of this change in assumptions is added to the top level demand forecast scheme, and demand probability distributions used to compute the displayed forecasts are adjusted accordingly [**Slide 13**].
- Next, Kevin decides to examine how the forecast for the Rome product group is computed from the various elements of forecast data that are resident in AX. [**Slide 14**] This brings up a more detailed view [**Slide 15**] which shows forecast data following from several perspectives, including (1) extrapolation for recent ordering trends, (2) projections based on historical ordering patterns in previous years of the current customer clusters, (3) the consolidated forecast built from the forecast data that was provided by various sales staff members, and (4) the set of actual confirmed orders for the forecast periods. The weights listed on the right indicate the relative influence of each of these data sources in producing the final forecast, and the result is visualized at the bottom of the display. It is possible for Kevin to adjust this influence by changing the weights listed on the right. In this case, however, Kevin is happy with the distribution of weights however he nonetheless feels that the consolidated forecast for the 2nd period is a bit too high (this intuition is based on information that Kevin has that is outside the system). He drags the point estimate down to finalize the forecast [**Slides 16-19**], and this pops him back to the top level screen [**Slide 20**]. (**note: slides 18-19 should be deleted in my opinion**).
 - The adjustment to the Rome product group is annotated as a 2nd changed assumption on the top level screen [**Slide 20**]. At this point Kevin is satisfied with the forecast and clicks ok to finish [**Slide 21**]. The forecast is stored in AX for subsequent use within the Supply Planning role center.

Next task: *Determine Rough Cut Supply Plan* task.

Task 2: Determine Rough Cut Supply Plan

Task Goal: Given the current product inventory, the current demand forecast, and supplier lead time and capacity constraints, determine

- the quantity of each product that needs to be ordered from suppliers this period
- an allocation of products to customer demand over the next n periods

that minimizes inventory cost (maximizes profit) and maximizes customer service level (% of orders received in the period in which demand is projected).

Persona: Eduardo (production planner), Ellen (warehouse manager).

Trigger: Regularly, each period, once the task “Produce customer demand forecast” is done. For purposes of the demo we will assume once a month.

Input Data:

- A demand forecast for all product and customer groups (expressed as probability distributions), together optionally with a target customer service level constraint for the firm's gold customers.
- Current inventory, including products currently in the warehouse, and those associated with open purchase orders that are expected to arrive in future periods (according to lead time constraints)

Outcome: A quantity for each product group that must be ordered from suppliers in the current period to meet projected demand in future periods at the final determined service level.

Task Steps:

- Eduardo and Ellen meet face-to-face and load up the current demand forecast produced and communicated by Kevin into the Supply Planning Tool. The supply Planning tool is first run to compute the optimal allocation for the current baseline weighting of inventory cost minimization and customer service level satisfaction. This solution is displayed in the top-level supply planning screen [**Slide 23**].
- Eduardo and Ellen first notice that the target service level constraints for gold customers is not being met (shown in the histogram on the right of the display). In response they first explore the global consequences of placing greater emphasis on meeting a high service level (at the potential expense of increased inventory and production costs). [**Slides 31-35**] (**note: These slides should go before the individual drill down into constituent customer group service levels**). After the supply plan is recomputed, they confirm that the service level constraint is now being satisfied.
- To better examine the tradeoff that has been achieved, they click on the customer group histogram [**Slide 24**] to bring up a more detailed view of the impact of the current global goal mix setting on projected inventory cost and profit. [**Slide 25**]. This display shows the expected inventory cost, profit and product turnover for each customer group for each of a number of service levels. (this is derived by computing a solution with a number of global goal mix settings and interpolating). If a given customer group (e.g., the large gold customer group) has an associated service level constraint, this would appear as a bar that cannot be violated. The draggable icon indicates the service level achieved for that customer group for the current rough cut supply plan.
- Eduardo and Ellen notice that the service level associated with the large Silver

customer group (which does not have a service level constraint) is unacceptably low in the current plan. By dragging the service level icon left or right [**Slide 26-28**], the plan is adjusted on a customer group specific level. As the service level of Large Silver is increased, the plan is recomputed and a corresponding decrease in the service levels achieved for other customer groups may be observed. In this case, Eduardo and Ellen are happy with the adjusted plan that better attends to the Large Silver customer group, and they click ok make this change permanent and return to the top level screen. [**Slide 29**].

- At this point, Eduardo and Ellen turn attention to the unfilled components of the current demand forecast. For each unfilled component of the current demand forecast, an analysis is made of various possibilities for meeting this unfulfilled demand. They look both for (1) excess manufacturing capacity for other product groups that could be re-purposed to lessen expected shortfalls, and (2) additional supplier capacity that could be obtained to increase supply. In this case, they are able to further improve the supply plan by repurposing manufacturing capacity. (In the underlying model this would amount to adjusting one or more capacity constraints and rerunning it). [**This step is not currently depicted in the slides**].
- The final supply plan to be adopted is then saved [**Slide 36-38**] and the quantities of each product group that must be ordered in the current period are generated. This supplier order data is then fed into a purchase requisition worksheet and forwarded to the Purchasing Department, where purchasers will then group demands for different products with different vendors to negotiate a discount (or consolidate with a single vendor to reach the quantity agreed in a frame agreement that triggers a large discount).

Next Task: *Generate Supplier Purchase Orders* task.

5. Differences with respect to Current ERP functionality

The optimization models and decision-support capabilities that we have prototyped and advocated in this research differ rather significantly from the tools and functionality provided in current ERP systems. We can make the following major distinctions:

- *Spreadsheet-like computational tools* – The constraint-based models we have promoted are very open. Most anything in the model can be an input parameter or decision variable, and the user can play a broad range of what-if analyses by fixing various parameters. For example, our basic Supply Planning model computes a reordering plan than minimizes inventory and delay costs (balancing customer service level achieved against inventory holding costs). If the user instead fixes customer service level at 100% as an input parameter, the model will compute the consequences with respect to increased inventory cost (assuming the base model does not achieve 100% service level). In contrast, existing ERP production plans are generally designed to take a fixed set of input parameters and compute a fixed set of outputs.
- *Interactive Analysis* - From a user experience point of view, our model computes results in real-time and is designed to support interactive analysis and problem solving. In contrast, the standard methodology of using existing ERP production planning tools in a supply planning context tends toward batch runs (potentially for extended periods) to compute results, and leads to much less analysis of options. Moreover, since there is considerable uncertainty associated with longer-term demand forecasts, the use of detailed deterministic methods (where “ghost orders” are created to enable computations, etc.) is somewhat of a mismatch here. Why plan in full detail for a series of uncertain events that are most certainly going to play out differently?
- *Most probable outcome reasoning* – Our models compute and exploit probability distributions that capture the uncertainty associated with demand forecasts to compute results that reflect the “most probable outcome”. The same concepts can be straightforwardly applied to the modeling of other uncertain constraints, such as supplier lead times and manufacturing yields. Existing ERP production planning tools utilize deterministic “expected” values (e.g., standard lead times, “pseudo orders”) and produce less accurate results.
- *Risk Assessment* – Utilizing our techniques for extracting probability distributions from historical information, our baseline supply planning model produces a risk-neutral result. However, by appropriately skewing the underlying distribution information, the model can be used to compute the consequences of more risk-taking or risk-averse decision-making and examine tradeoffs. Existing ERP tools do not provide this sort of analysis capability.
- *Visualization* – We propose the use of visualizations of the production and supply network process, linked to underlying business processes, to characterize the expected performance of a given solution and provide an intuitive, high-level basis for using optimization tools to evaluate decision alternatives. Existing interfaces to ERP systems provide very little support for structuring of business

processes and analysis of business decisions. We envision the introduction of a higher-level business process representation layer as a super-structure for inserting advanced analytic and optimization capabilities in a way that directly promotes and enhances role-based decision-making.

6.0 Conclusions and Next Steps

This research has taken initial steps toward the development and use of advanced optimization models within the demand and supply planning processes of manufacturing organizations. Our approach combines several key ideas. First, we emphasize the use of modern constraint programming and reasoning techniques as a basis for constructing optimization models; the flexibility and scalability of these techniques allow the construction of powerful “spreadsheet style” decision-support tools, where different model parameters, variables and constraints can be selectively manipulated and the model can be repeatedly solved in real time to carry out sophisticated what-if analyses interactively. Second, we emphasize the development of optimization models that explicitly take the uncertainty associated with demand and manufacturing processes into account. Such models provide much more accurate results than the “expected-value” reasoning that underlies current ERP planning modules, and hence a stronger basis for decision-support. Further, we have shown that this does not preclude the use of deterministic solving techniques; in particular, the use of scenario-based reasoning to embed a deterministic solver in a larger stochastic optimization procedure provides a flexible basis for taking domain uncertainty into account. Third, we emphasize direct integration of optimization models as additional decision support tools provided by an ERP system, and as such, these models must draw their inputs directly from ERP data. Since ERP data does not explicitly represent the uncertainty associated with data values, one challenge is extracting this information from the historical data that is resident in the ERP system. Our research has demonstrated that the use of clustering and kernel density estimation techniques from the field of machine learning can provide a means. Finally, we emphasize a coupling of business process modeling with Microsoft’s concept of role-based decision-making as a framework for embedding advanced analytic decision support capabilities into an ERP system. We envision a higher-level ERP representation layer that directly associates advanced analytics and optimization models with the decision points that they are designed to support, and enables visualization and propagation of the non-local impacts of decisions across role boundaries.

There are several natural next steps to furthering this research. One necessary next step is to consider the demand and supply planning data and problems of a larger set of manufacturing organizations. The supply planning tool formulated in this initial effort has been influenced by Hoshino USA’s supply planning problem, and this influence has led to a model that emphasizes its role as a distributor. The consideration of supply planning environments that additionally require that substantial internal production processes be taken into account is necessary to refine and broaden the scope of the prototype model. A second next step is to consider reformulating the prototype model as a constraint programming (CP) model. As the model is generalized to account for the scheduling constraints typically associated with managing internal production processes, the advantages of a CP model in efficiently solving these types of combinatorial problems is likely to argue for conversion of the current mathematical model. A final next step is to begin to consider the larger set of decisions that are made by users performing demand and supply planning roles, and identify additional role-based decision-support capabilities that optimization models could provide.

7.0 References

[Amram and Kulatilaka 1999] Martha Amram and Nalin Kulatilaka, *Real Options: Managing Strategic Investment in an Uncertain World*, Harvard Business School Press, 1999.

[MS 2007] “Microsoft Dynamics Role Tailored Business Productivity: Software Designed for your People”, Microsoft Corporation White Paper, March, 2007

[Savage et. al 2006] Sam Savage, Stefan Scholtes and Daniel Zweidler, “Probability Management”, *ORMS Today*, 33(1), February 2006.

[Tarim et. al 2006], S. Armagan Tarim, S. Manandhar and Toby Walsh “Stochastic Constraint Programming: A Scenario Based Approach”, *Constraints*, 11(1), 53-81, 2006.

[Van Hentenryck and Bent 2006] Pascal Van Hentenryck and Russell Bent, *Online Stochastic Combinatorial Optimization*, MIT Press, October, 2006.

A Supply Planning Model

In this appendix, we formally specify two variants of the model developed as the supply planning tool. The basic version (which was implemented in CPLEX and delivered as the core solver of the prototype supply planning tool) answers the question of what product quantities to order from suppliers when so as to minimize inventory and customer service delay costs (or equivalently to maximize profits). The extended version answers the additional question of what products and customers to focus on, given the availability of some amount of additional capital to conduct a marketing campaign. We consider each model individually in the subsections below. Then we summarize the overall stochastic optimization framework and specify how the set of k generated solutions are combined to produce the final model output.

A.1 Determining optimal reordering decisions

For purposes of formulating the basic supply planning model, we first introduce the following basic notation:

- $t \rightarrow$ denotes a time epoch and $0 < t < H$, where H is time horizon.
- $c \rightarrow$ denotes a customer and belongs to the set C .
- $p \rightarrow$ denotes a product and belongs to the set P .
- $s \rightarrow$ denotes a supplier and belongs to the set S .

Now, let us define the following variables and parameters:

- $d_t^{c,p} \rightarrow$ The demand from customer c for product p to be serviced by time period, t ($k \leq t \leq H$).
- $o_t^{p,s} \rightarrow$ The quantity of product, p ordered by the distributor from supplier s at time period, t ($0 \leq t \leq H$).
- $r_t^{p,s} \rightarrow$ The quantity of product, p received by the distributor from supplier s at time period, t . From the total quantity of product p , i.e. $\sum_s r_t^{p,s}$,
- $\alpha_{t,t_1}^{c,p} \rightarrow$ the quantity of product p allocated to service customer c 's demand for p in time period t_1 .
- $i_t^p \rightarrow$ The quantity of product p in inventory at time period t .

- $\gamma^c \rightarrow$ the service level penalty for delaying any product delivery to customer, c for one time period.
- $q_t^{p,s} \rightarrow$ The quantity of product p delivered by supplier s after t time periods.
- $l_{p,s} \rightarrow$ The lead time for delivery of product p from supplier s .
- $cap^{p,s} \rightarrow$ The maximum quantity of product p that supplier s can produce per period.
- $\psi^{c,p} \rightarrow$ The sales price of product p to customer c .
- $\phi^{p,s} \rightarrow$ The cost of product p from supplier s .
- $ContCap \rightarrow$ The capacity of a container used to transport products from a supplier to the distributor.
- $nbrCont_t^s \rightarrow$ The number of containers arriving from supplier s in time period t . We assume that it takes negligible amount of time to ship a container (relative to the duration of a time period).
- $InvCost \rightarrow$ The unit cost of holding a given product in the distributor's warehouse
- $ShipCost \rightarrow$ The cost of shipping a container full of products

Given these definitions we can specify the base optimization problem of interest as follows:

Variables : $\alpha_{t,t_1}^{c,p}, o_t^{p,s}, nbrCont_t^s$

Inputs : $d_t^{c,p}, q_t^{p,s}, i_t^p, l_{p,s}, ContCap, InvCost, ShipCost, \psi^{c,p}, \phi^{s,p}$

$$\begin{aligned}
 \text{Maximize :} \quad & \sum_{c,t,p,t_1 \leq t} \alpha_{t,t_1}^{c,p} * \{\psi^{c,p} - \gamma^c * (t - t_1)\} - \sum_{t,p,s} r_t^{p,s} * \phi^{p,s} \\
 & - \sum_{t,p} i_t^p * InvCost - nbrCont_t^s * ShipCost \quad (1)
 \end{aligned}$$

$$\begin{aligned}
 \text{s.t.} \quad & \forall t_1, c, p : \sum_t \alpha_{t,t_1}^{c,p} = d_{t_1}^{c,p} \\
 & \forall t, p : \sum_s r_t^{p,s} + i_{t-1}^p - \sum_{c,t_1} \alpha_{t,t_1}^{c,p} = i_t^p, t_1 \leq t \\
 & \forall t, s, p : r_t^{p,s} \leq cap^{p,s} \\
 & \forall t, s, p : r_t^{p,s} \leq o_{t-l_{p,s}}^{p,s} \\
 & \forall p, s : l_{p,s} = \max_i l_{p_i,s_i} \\
 & \forall t, s : \sum_p r_t^{p,s} = ContCap * nbrCont_t^s
 \end{aligned}$$

The objective specified in equation 1 is to maximize the net profit. The first term specifies the total profit obtained from delivered products, factoring in cost penalties associated with late deliveries. The second term specifies the total costs incurred for products received from suppliers. The third term specifies the total inventory holding cost incurred, factoring out the products that have just been received from suppliers.

The maximization of this objective is subject to five constraints:

1. The delivered amount of any product must be less than or equal to the demand amount (Constraint 1)
2. Any product excess in a given time period is due to supplier demand (Constraints 2)
3. The quantity of products received in a given period must be less than or equal to the supplier's capacity (Constraint 3)

4. The quantity of products received in a given period must be less than or equal to the amount ordered (Constraint 4)
5. The lead time of a given product order is the max of the lead times of each supplier. (Constraint 5) - This addresses situations when a combination of suppliers $\times s_i$) delivers sub-products p_i that are assembled into a final product p
6. The quantity of products received in a given period must equal the number of containers shipped (Constraint 6).

A.2 Determining where to put marketing effort

To address the extended optimization problem of determining which customers and products to focus on if extra capital is available to support a marketing campaign, we introduce two additional parameters:

- $v_t^p \rightarrow$. The value of product p to the customer at time period t .
- $\delta_t^{c:p}(e_t^p) \rightarrow$ A function indicating the worth e_t^p of increased demand from customer c for product p at time period t due to advertising.

We also remove consideration of supplier batching (container) constraints to simplify the presentation (it is straightforward to add these constraints back into the extended model). Given this assumption, the extended optimization problem is formulated as follows:

$$\begin{aligned}
\text{Variables :} & \quad e_t^p, o_t^{p,s} \\
\text{Inputs :} & \quad e, d_t^{c,p}, \delta_t^{c,p}(), q_t^{p,s} \\
\max & \quad \sum_{c,t,p,t_1 \leq t} \alpha_{t,t_1}^{c,p} * \{v_t^p - \gamma^c * (t - t_1)\} - \sum_{t,p} i_t^p * \text{cost} \quad (2) \\
\text{s.t.} & \quad \forall t, c, p : \sum_{t_1} \alpha_{t,t_1}^{c,p} \leq d_t^{c,p} + \delta_t^{c,p}(e_t^p) \\
& \quad \forall t : \sum_s r_t^{p,s} - \sum_{c,t_1} \alpha_{t,t_1}^{c,p} = x_t^p \\
& \quad \forall t, p : i_t^p = i_{t-1}^p + x_t^p \\
& \quad \forall t, s, p : r_t^{p,s} \leq \sum_{t' < t} q_{t-t'}^{p,s} \\
& \quad \forall t, s, p : r_t^{p,s} \leq \sum_{t' < t} \{o_{t'}^{p,s} - r_{t'}^{p,s}\} \\
& \quad \forall t, p : \sum_{t,p} e_t^p \leq e
\end{aligned}$$

The above formulation differs from the basic formulation only in the objective function (equation 2) and constraints 1 and 5. In the objective function we replace $\psi^{c,p}$ with v_t^p . Constraint 1 captures the increased demand due to the extra investment on each product. Constraint 5 accounts for the fact that sum of investments on all products must be less than or equal to the overall investment.

A.3 Combining results from several scenarios

Recall that we employ a scenario based approach to solve the models provided in the earlier section, primarily to account for the fact that there is uncertainty associated the input data (demand distributions). Specifically, we generate k samples from the underlying demand distribution, and obtain solutions for each sample. Note that for a fixed sample of demands, we need only to solve a linear program. After obtaining solutions for each generated sample (scenario), the final solution is obtained by combining those individual solutions. For problem 1, this corresponds to: (a) obtaining the quantity to be ordered at a specific time step for each sample; and (b) computing the

expected value ($= \sum_{i < n} a_t^{p,s}(i) * p(i)$, where $p(i)$ is the probability associated with that specific sample of demand).

B. User Interface Design Sketches for Task Scenarios

CMU Project

UI Sketches for scenarios

Slide 1

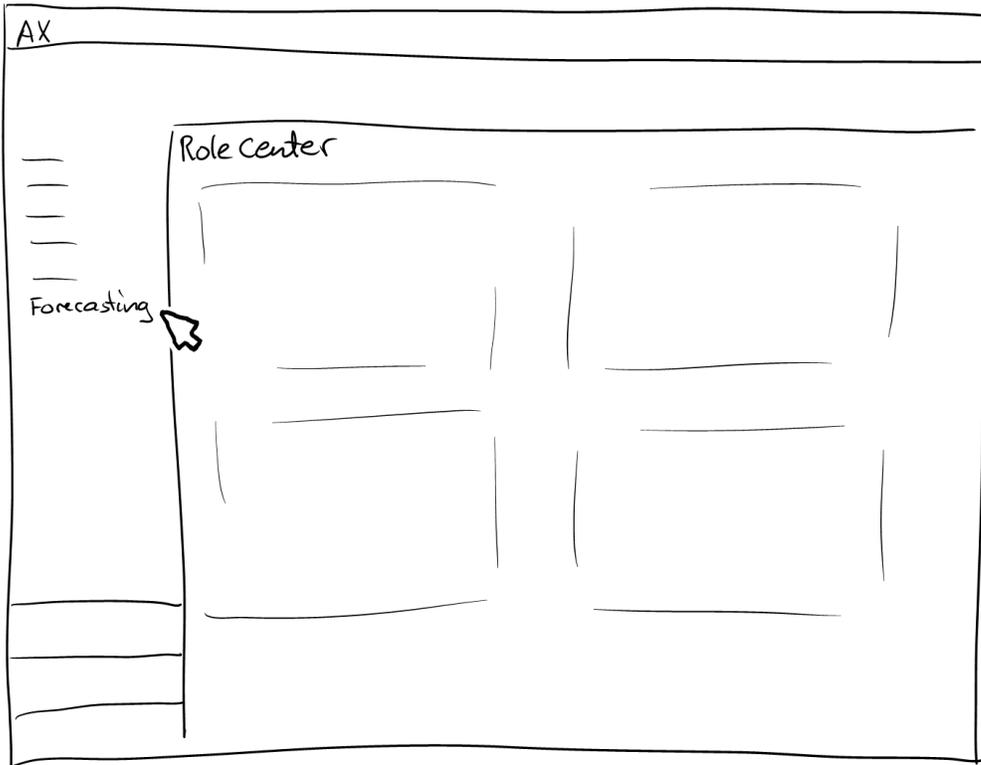
Produce Demand Forecast

User: Kevin, sales manager

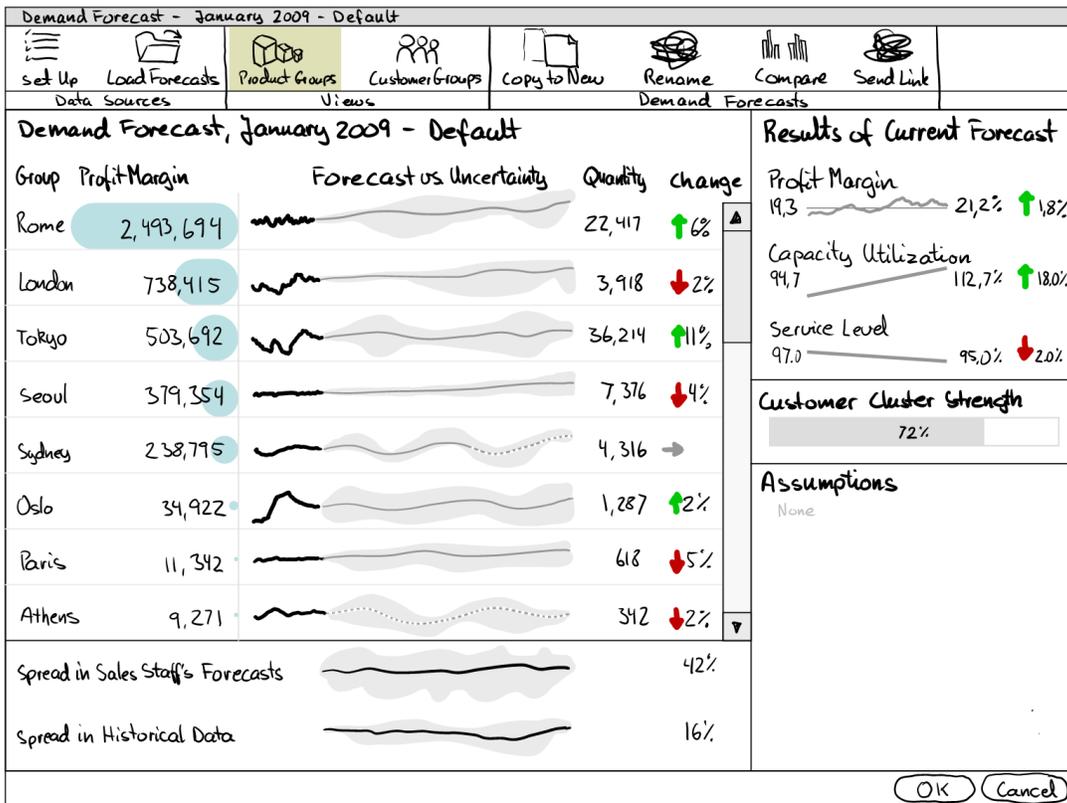
Slide 2



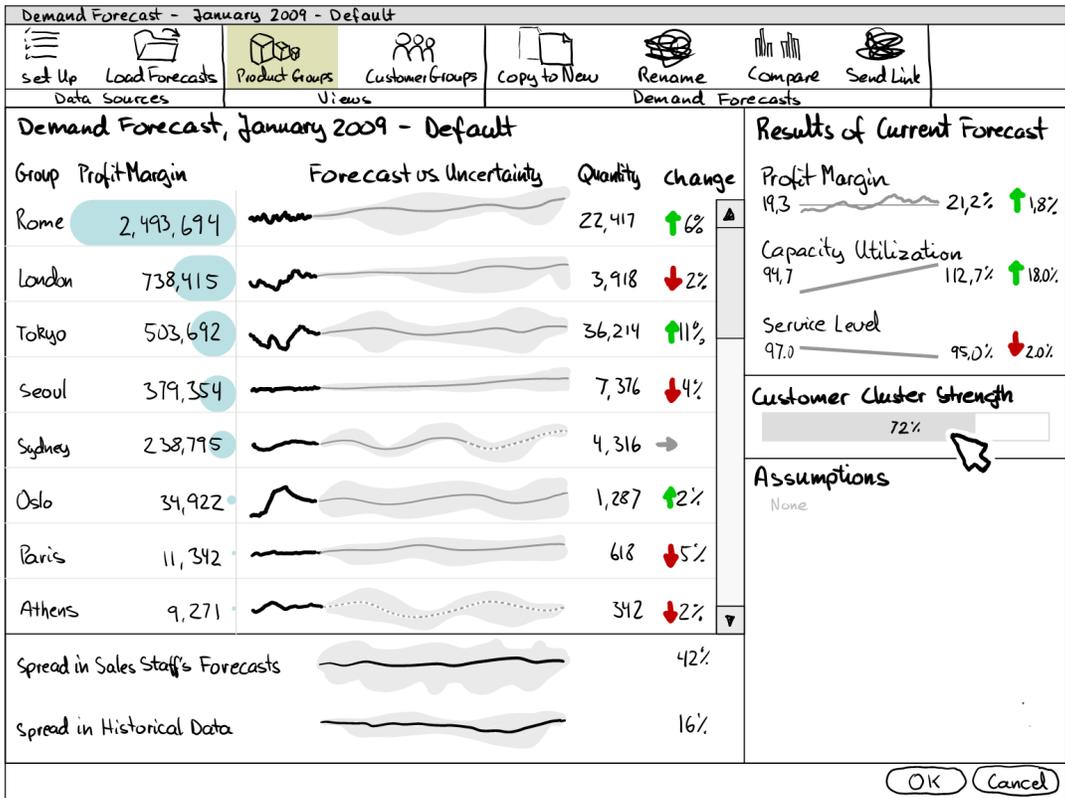
Slide 3



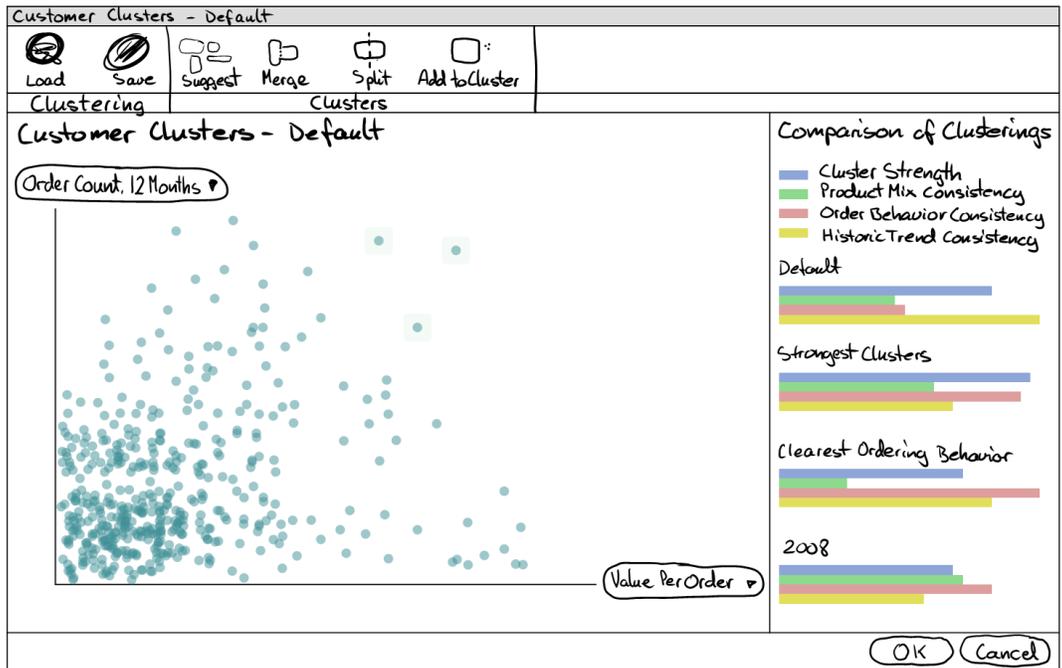
Slide 4



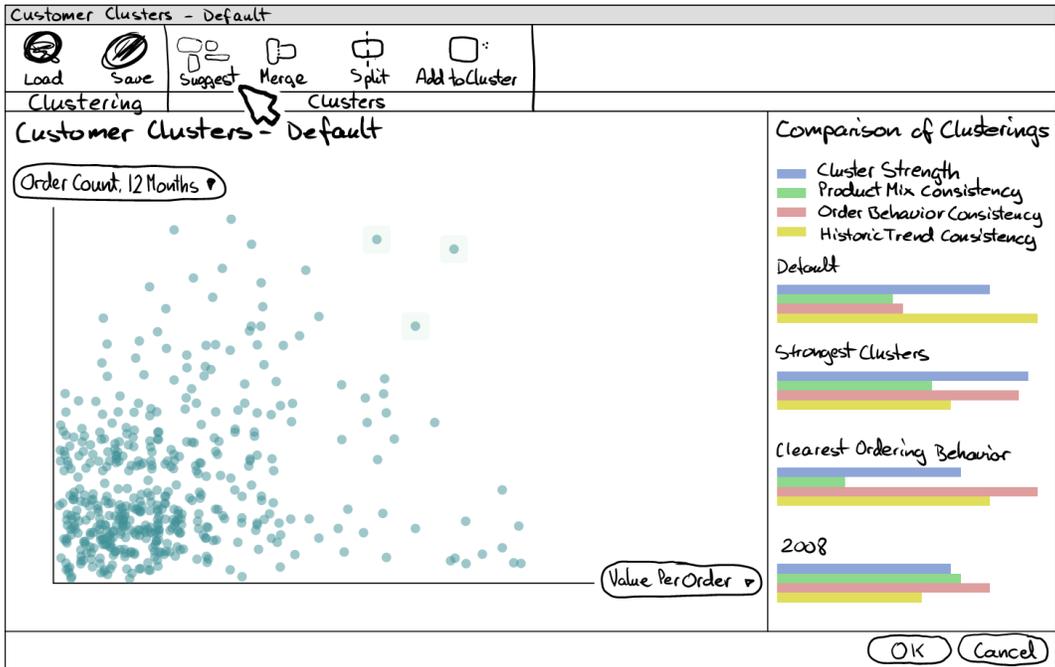
Slide 5



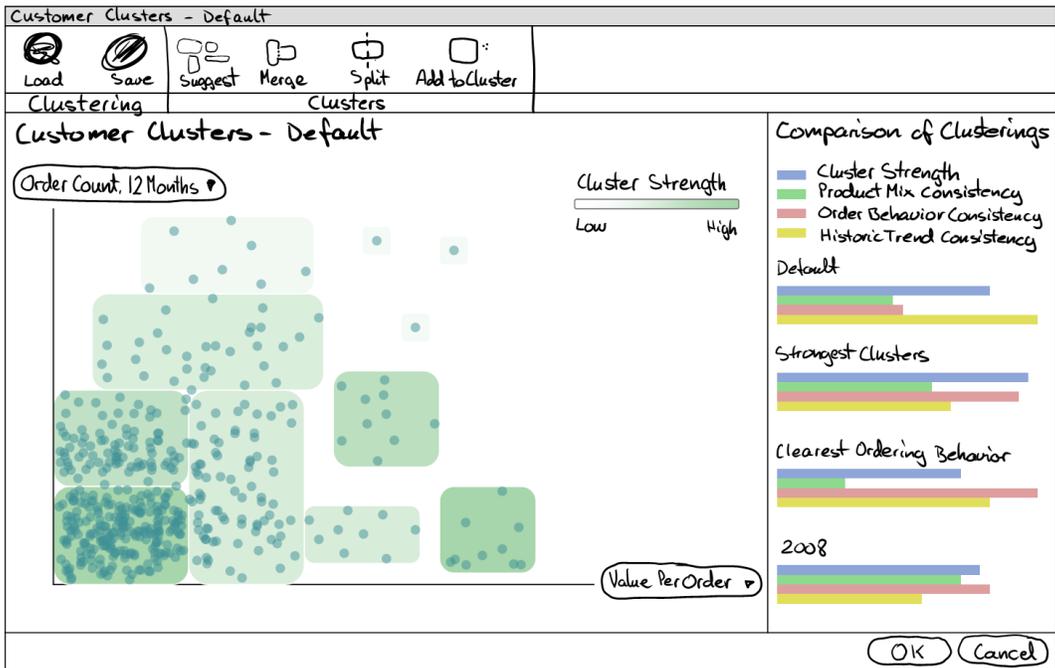
Slide 6



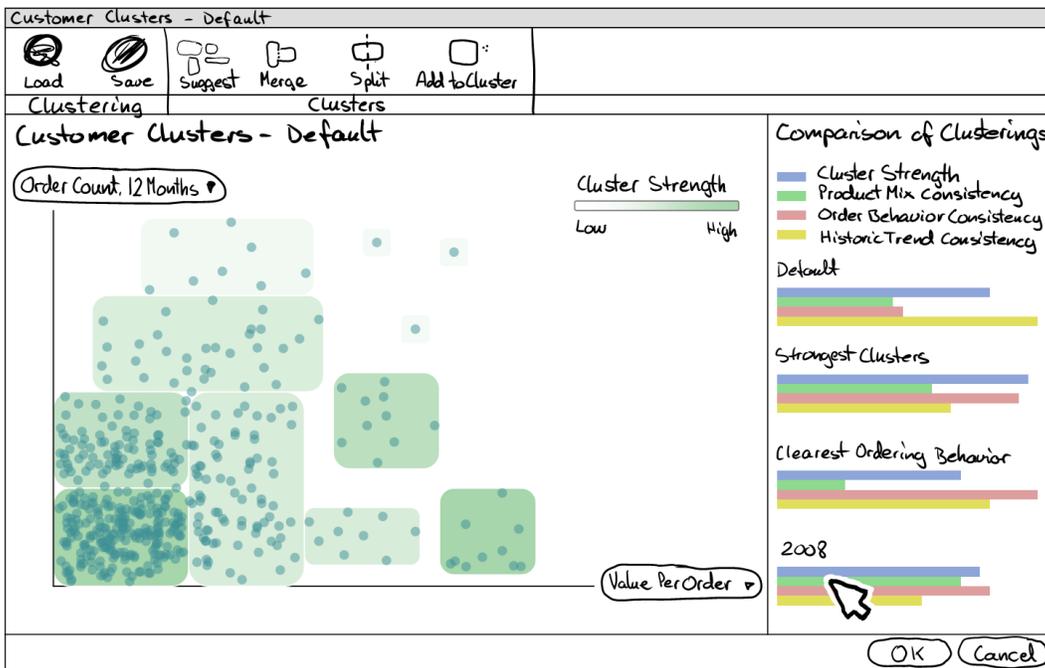
Slide 7



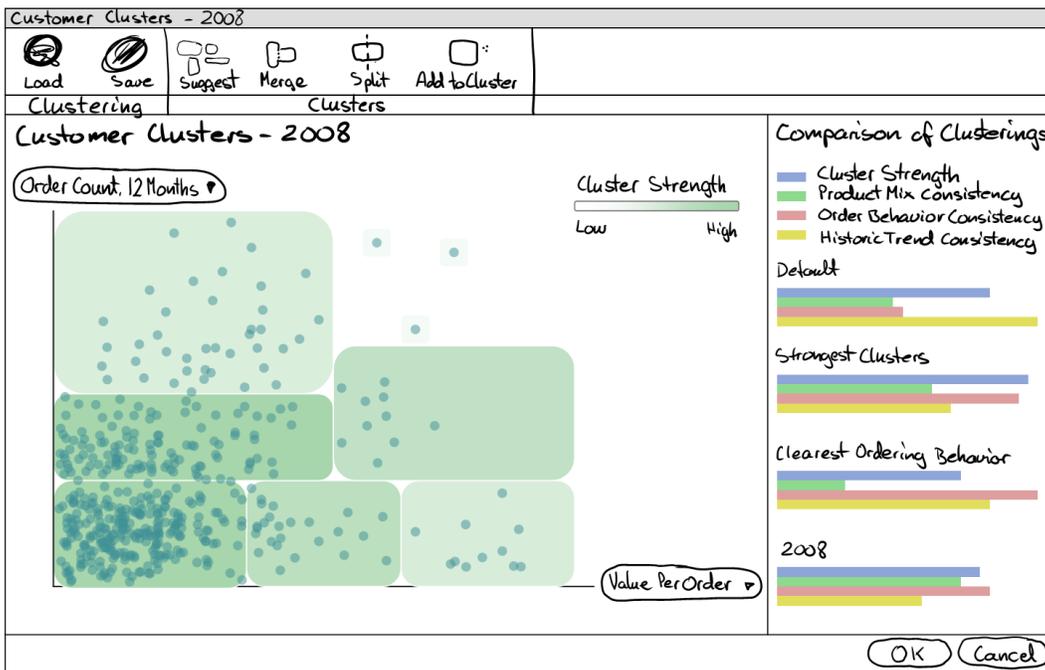
Slide 8



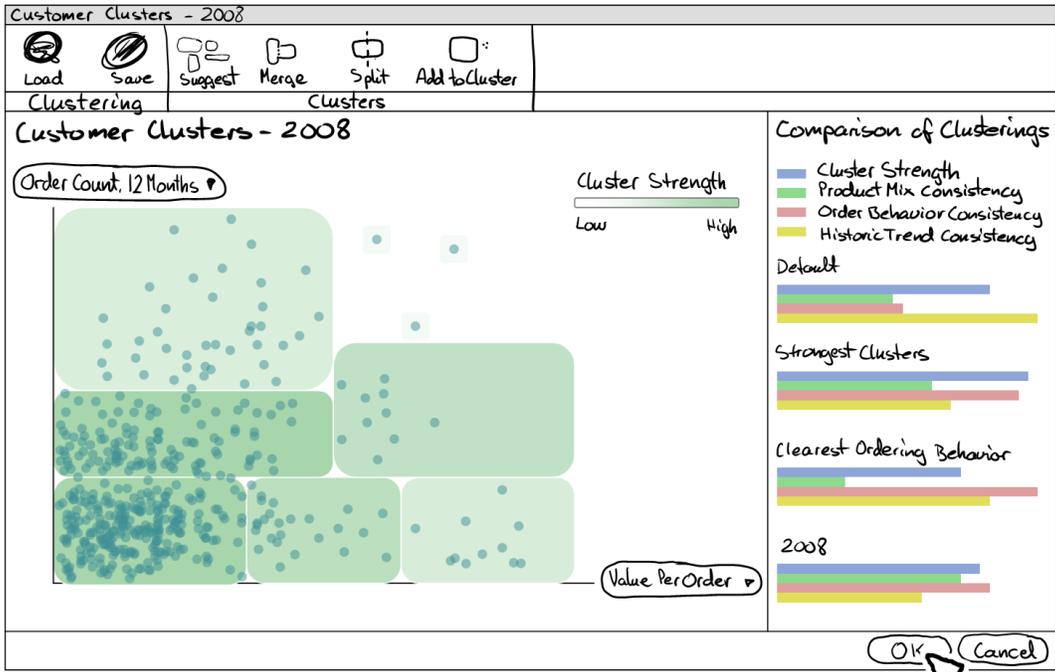
Slide 9



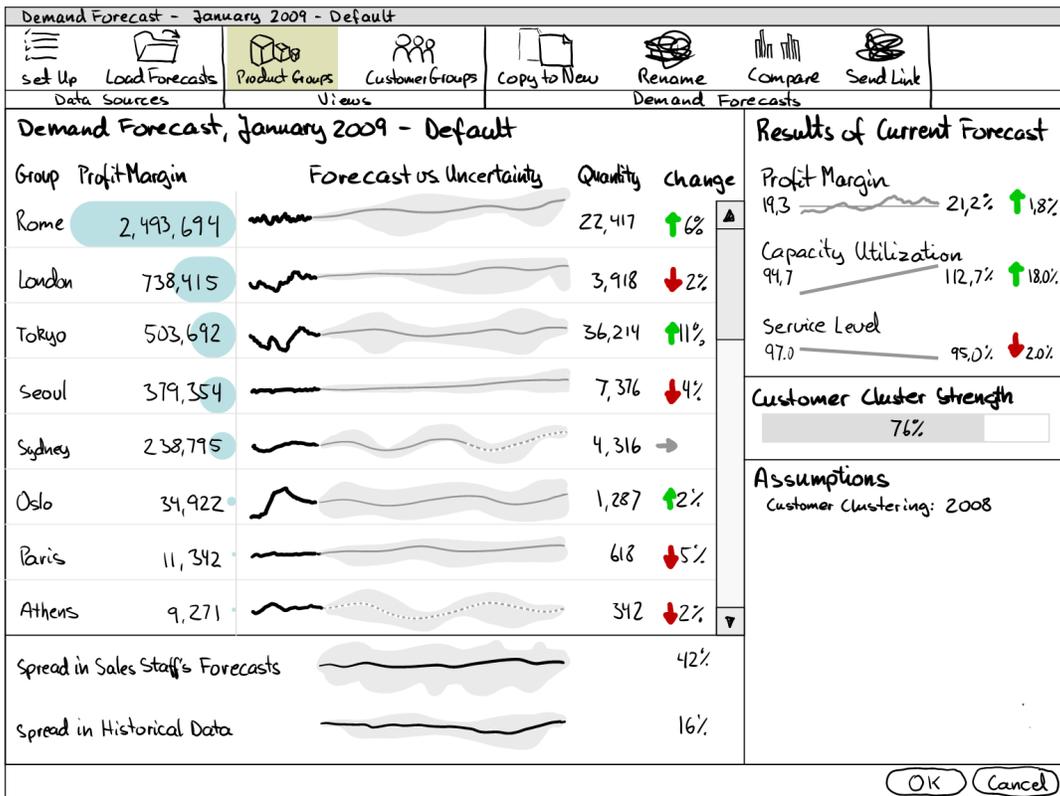
Slide 10



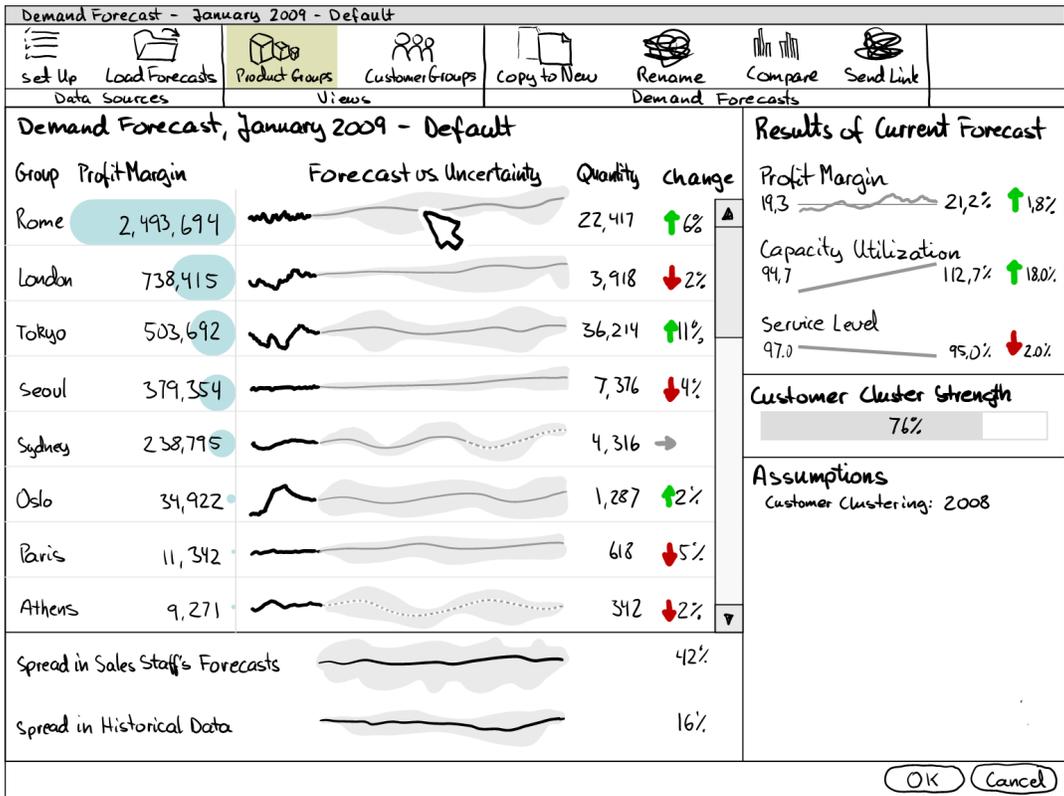
Slide 11



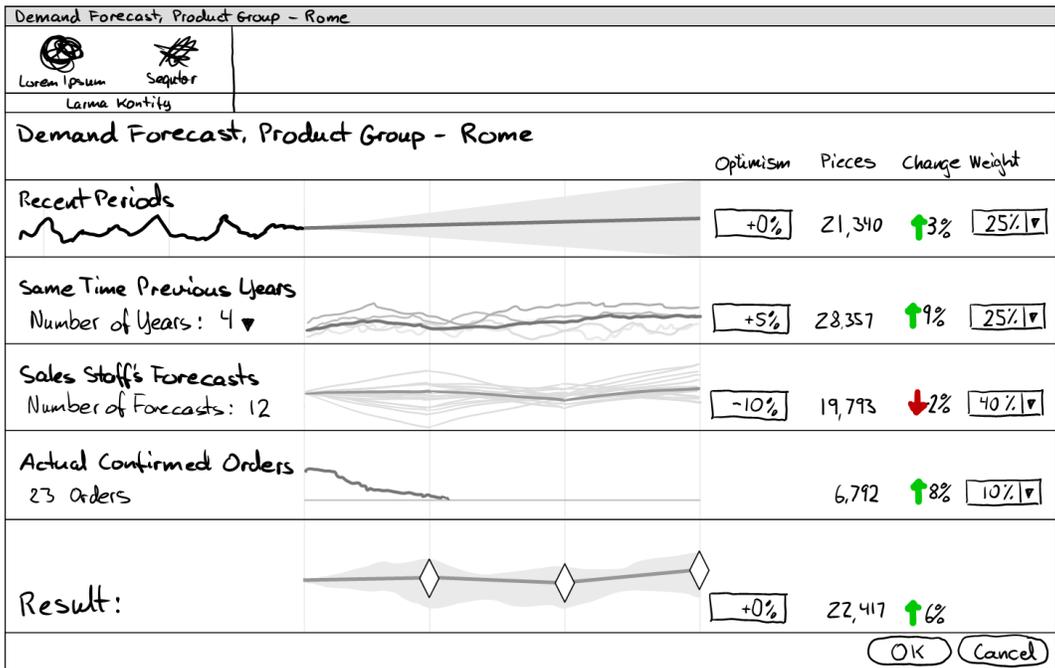
Slide 12



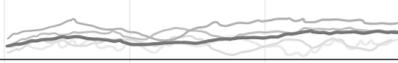
Slide 13



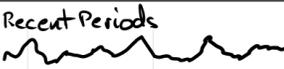
Slide 14



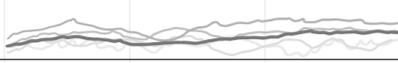
Slide 15

Demand Forecast, Product Group - Rome				
 Lorem Ipsum	 Sagutor			
Lanna KontiGy				
Demand Forecast, Product Group - Rome				
		Optimism	Pieces	Change Weight
Recent Periods 		<input type="text" value="+0%"/>	21,340	<input type="text" value="↑3%"/> <input type="text" value="25% ▼"/>
Same Time Previous Years Number of Years: 4 ▼		<input type="text" value="+5%"/>	28,357	<input type="text" value="↑9%"/> <input type="text" value="25% ▼"/>
Sales Staff's Forecasts Number of Forecasts: 12		<input type="text" value="-10%"/>	19,793	<input type="text" value="↓2%"/> <input type="text" value="40% ▼"/>
Actual Confirmed Orders 23 Orders			6,792	<input type="text" value="↑8%"/> <input type="text" value="10% ▼"/>
Result:		<input type="text" value="+0%"/>	22,417	<input type="text" value="↑6%"/>
		<input type="button" value="OK"/> <input type="button" value="Cancel"/>		

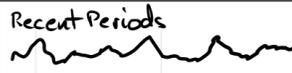
Slide 16

Demand Forecast, Product Group - Rome				
 Lorem Ipsum	 Sagutor			
Lanna KontiGy				
Demand Forecast, Product Group - Rome				
		Optimism	Pieces	Change Weight
Recent Periods 		<input type="text" value="+0%"/>	21,340	<input type="text" value="↑3%"/> <input type="text" value="25% ▼"/>
Same Time Previous Years Number of Years: 4 ▼		<input type="text" value="+5%"/>	28,357	<input type="text" value="↑9%"/> <input type="text" value="25% ▼"/>
Sales Staff's Forecasts Number of Forecasts: 12		<input type="text" value="-10%"/>	19,793	<input type="text" value="↓2%"/> <input type="text" value="40% ▼"/>
Actual Confirmed Orders 23 Orders			6,792	<input type="text" value="↑8%"/> <input type="text" value="10% ▼"/>
Result:		<input type="text" value="-2%"/>	21,968	<input type="text" value="↑4%"/>
		<input type="button" value="OK"/> <input type="button" value="Cancel"/>		

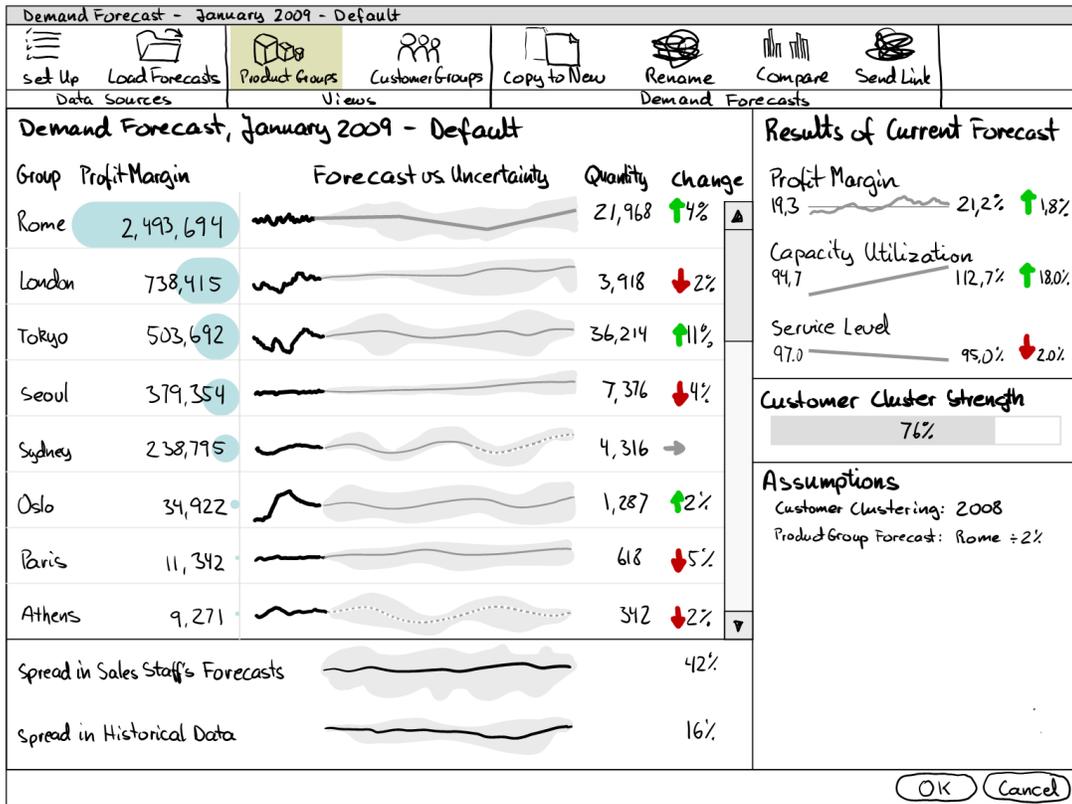
Slide 17

Demand Forecast, Product Group - Rome			
 Lorem Ipsum Lanna Kontig	 Sagutor		
Demand Forecast, Product Group - Rome			
	Optimism	Pieces	Change Weight
Recent Periods 	<input type="text" value="+0%"/>	21,340	<input type="text" value="↑3%"/> <input type="text" value="25% ▼"/>
Same Time Previous Years Number of Years: 4 ▼ 	<input type="text" value="+5%"/>	28,357	<input type="text" value="↑9%"/> <input type="text" value="25% ▼"/>
Sales Staff's Forecasts Number of Forecasts: 12 	<input type="text" value="-10%"/>	19,793	<input type="text" value="↓2%"/> <input type="text" value="40% ▼"/>
Actual Confirmed Orders 23 Orders 		6,792	<input type="text" value="↑8%"/> <input type="text" value="10% ▼"/>
Result: 	<input type="text" value="-2%"/>	21,968	<input type="text" value="↑4%"/>
<input type="button" value="OK"/> <input type="button" value="Cancel"/>			

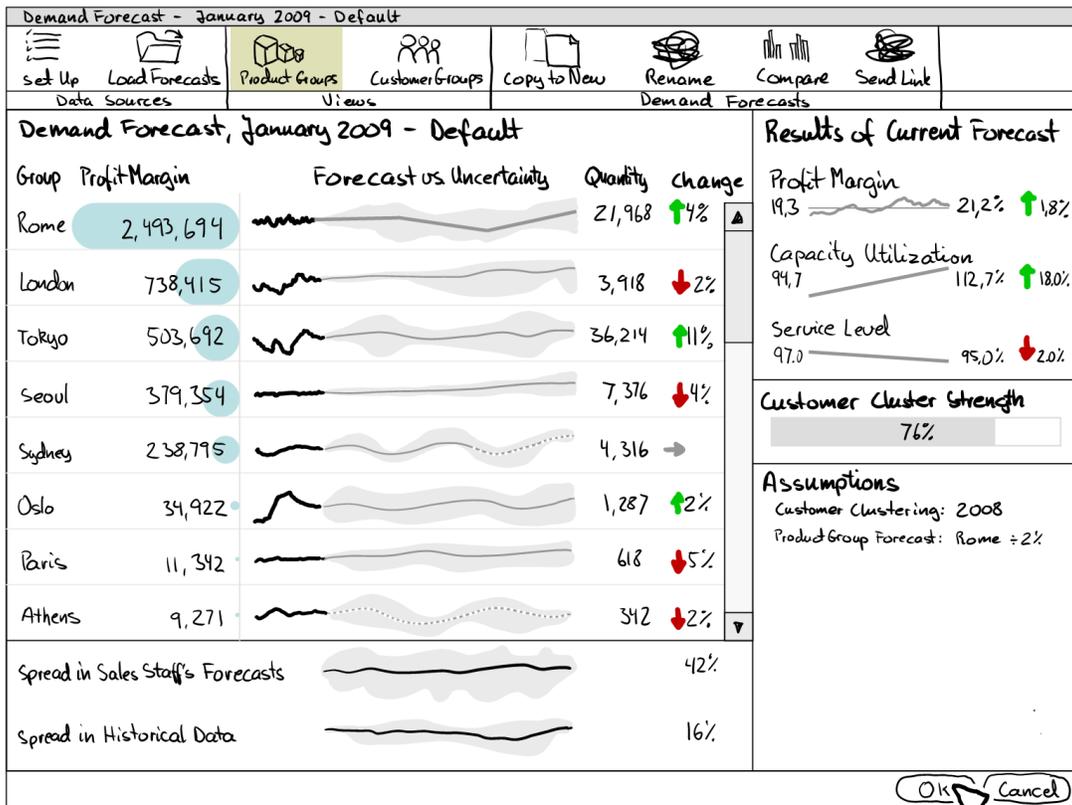
Slide 18

Demand Forecast, Product Group - Rome			
 Lorem Ipsum Lanna Kontig	 Sagutor		
Demand Forecast, Product Group - Rome			
	Optimism	Pieces	Change Weight
Recent Periods 	<input type="text" value="+0%"/>	21,340	<input type="text" value="↑3%"/> <input type="text" value="25% ▼"/>
Same Time Previous Years Number of Years: 4 ▼ 	<input type="text" value="+5%"/>	28,357	<input type="text" value="↑9%"/> <input type="text" value="25% ▼"/>
Sales Staff's Forecasts Number of Forecasts: 12 	<input type="text" value="-10%"/>	19,793	<input type="text" value="↓2%"/> <input type="text" value="40% ▼"/>
Actual Confirmed Orders 23 Orders 		6,792	<input type="text" value="↑8%"/> <input type="text" value="10% ▼"/>
Result: 	<input type="text" value="-2%"/>	21,968	<input type="text" value="↑4%"/>
<input type="button" value="OK"/> <input type="button" value="Cancel"/>			

Slide 19



Slide 20



Slide 21

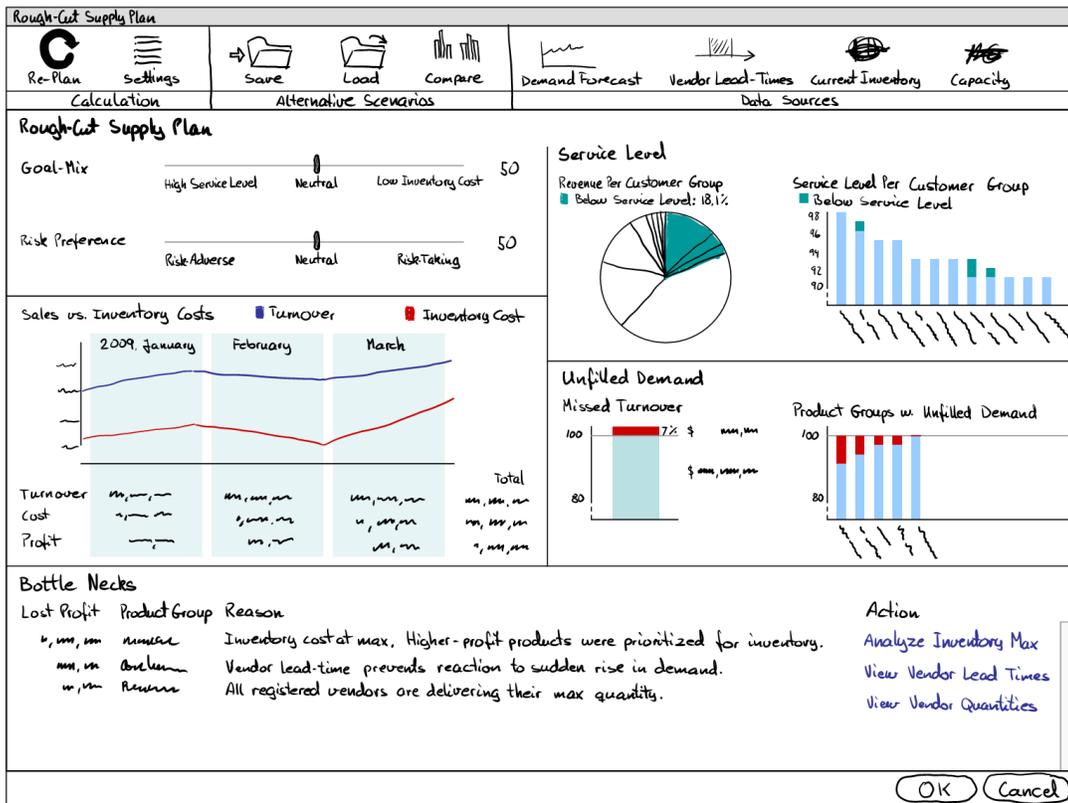
Generate Supply Plan

Users:

Eduardo (production planner)

Ellen (warehouse manager)

Slide 22



Slide 23

Rough-Cut Supply Plan

Re-Plan Settings Save Load Compare Demand Forecast Vendor Lead-Times Current Inventory Capacity

Calculation Alternative Scenarios Data Sources

Rough-Cut Supply Plan

Goal-Mix: High Service Level Neutral Low Inventory Cost 50

Risk Preference: Risk Averse Neutral Risk Taking 50

Sales vs. Inventory Costs

	2009 January	February	March	Total
Turnover	10,000,000	10,000,000	10,000,000	30,000,000
Cost	2,000,000	2,000,000	2,000,000	6,000,000
Profit	8,000,000	8,000,000	8,000,000	24,000,000

Service Level

Revenue per Customer Group: Below Service Level: 18.1%

Unfilled Demand

Missed Turnover: 7% \$ 1,000,000

Product Groups w. Unfilled Demand

Bottle Necks

Lost Profit	Product Group	Reason	Action
1,000,000	xxxxxx	Inventory cost at max. Higher-profit products were prioritized for inventory.	Analyze Inventory Max
100,000	xxxxxx	Vendor lead-time prevents reaction to sudden rise in demand.	View Vendor Lead Times
100,000	xxxxxx	All registered vendors are delivering their max quantity.	View Vendor Quantities

OK Cancel

Slide 24

Service Level Per Customer Group

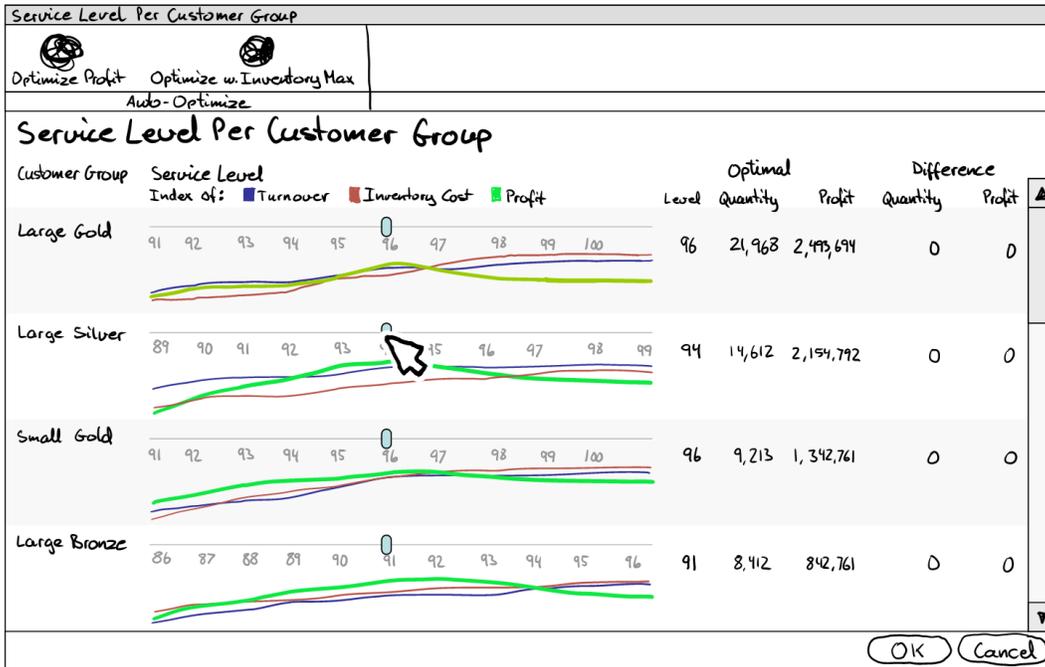
Optimize Profit Optimize w. Inventory Max Auto-Optimize

Service Level Per Customer Group

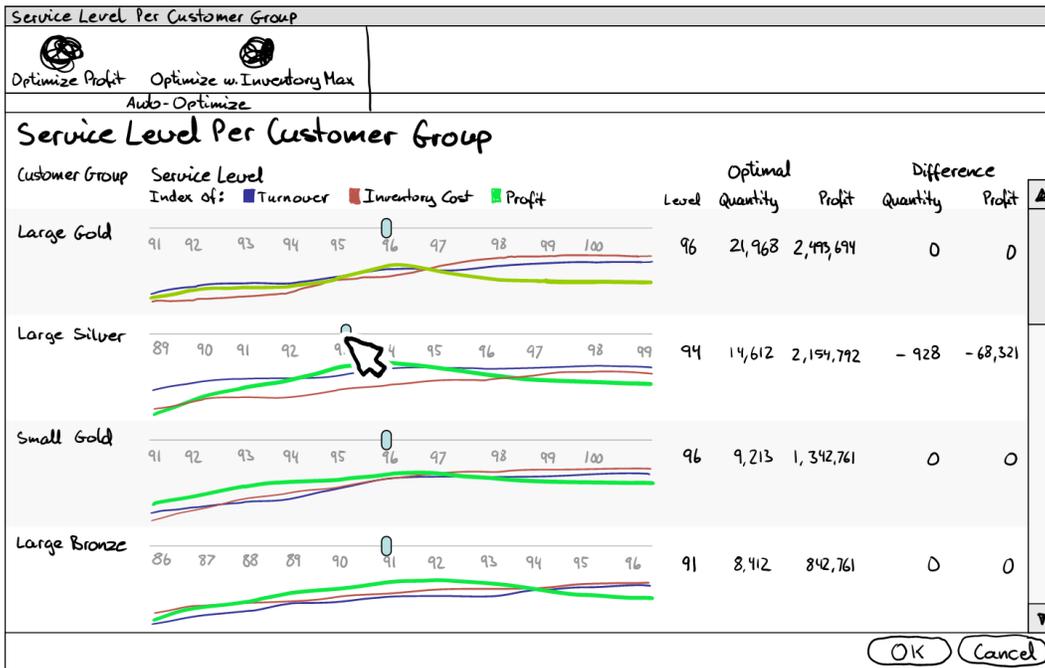
Customer Group	Service Level Index of: Turnover Inventory Cost Profit	Optimal		Difference	
		Level	Quantity	Quantity	Profit
Large Gold	91 92 93 94 95 96 97 98 99 100	96	21,963	2,493,694	0
Large Silver	89 90 91 92 93 94 95 96 97 98 99	94	14,612	2,154,792	0
Small Gold	91 92 93 94 95 96 97 98 99 100	96	9,213	1,342,761	0
Large Bronze	86 87 88 89 90 91 92 93 94 95 96	91	8,412	842,761	0

OK Cancel

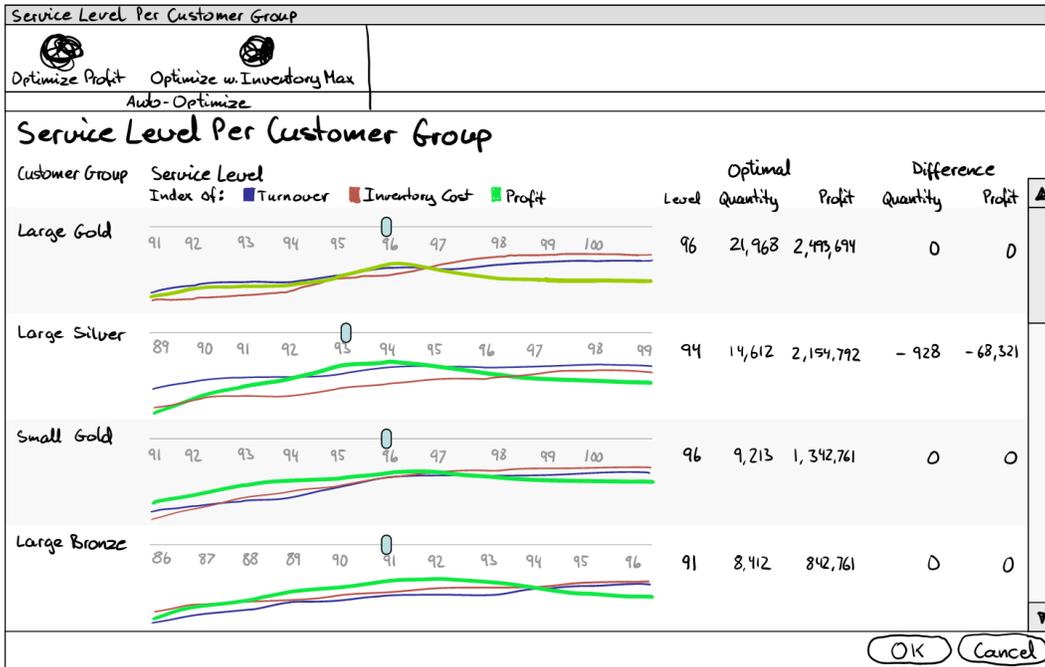
Slide 25



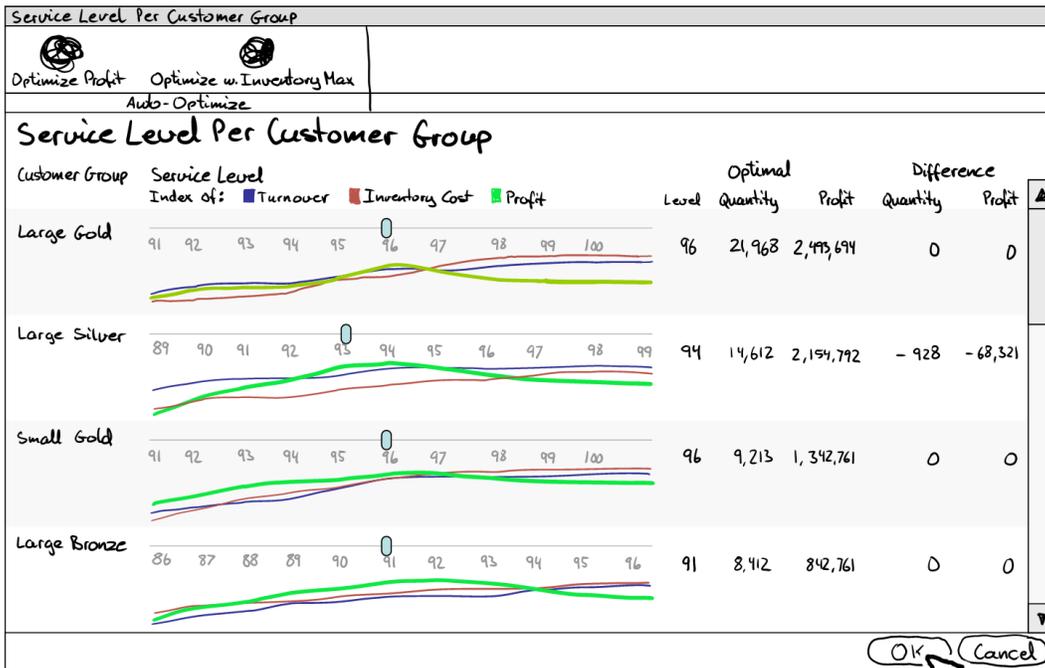
Slide 26



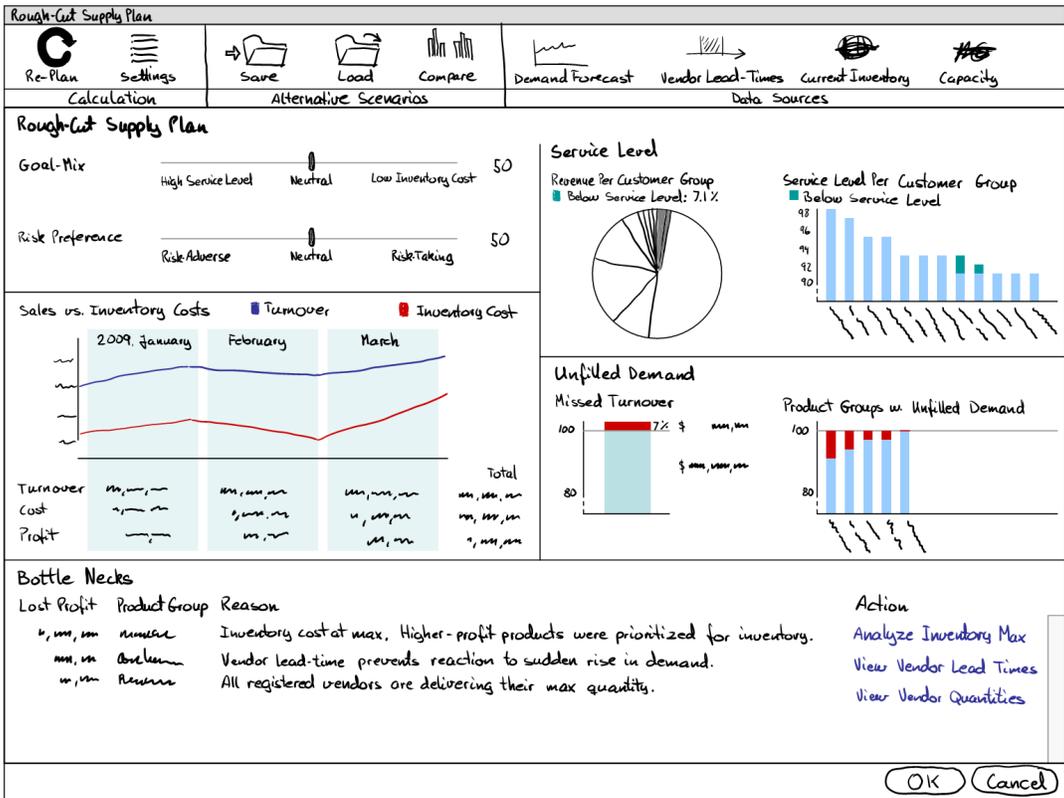
Slide 27



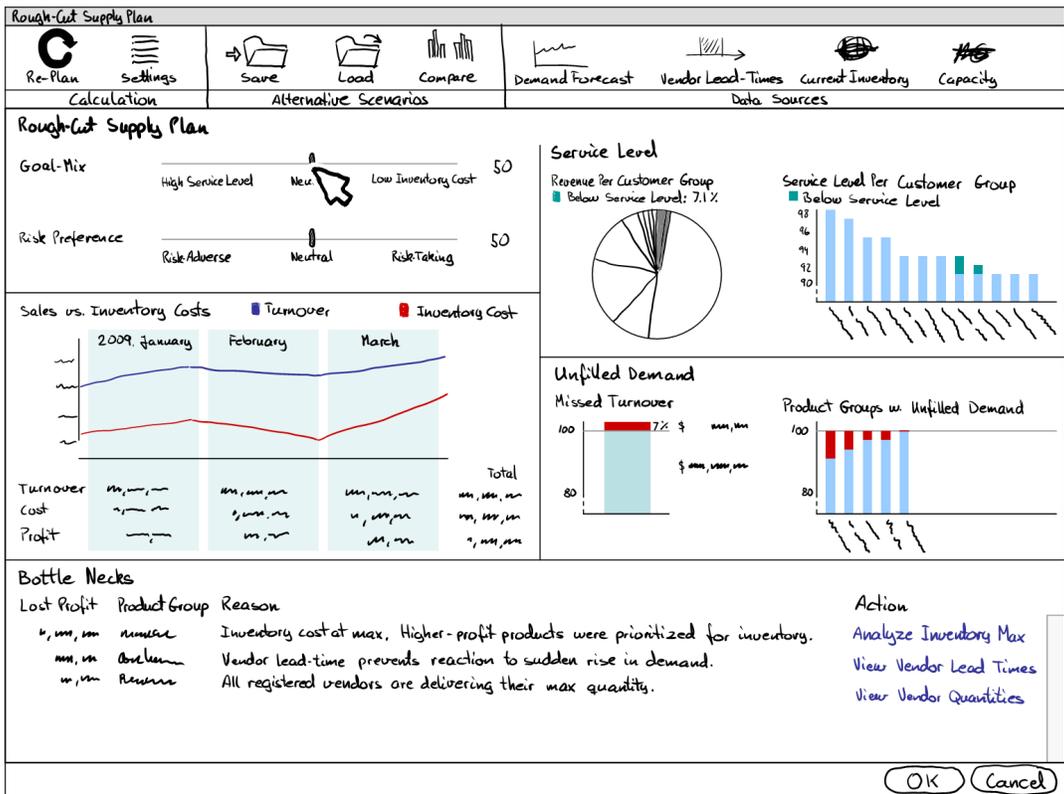
Slide 28



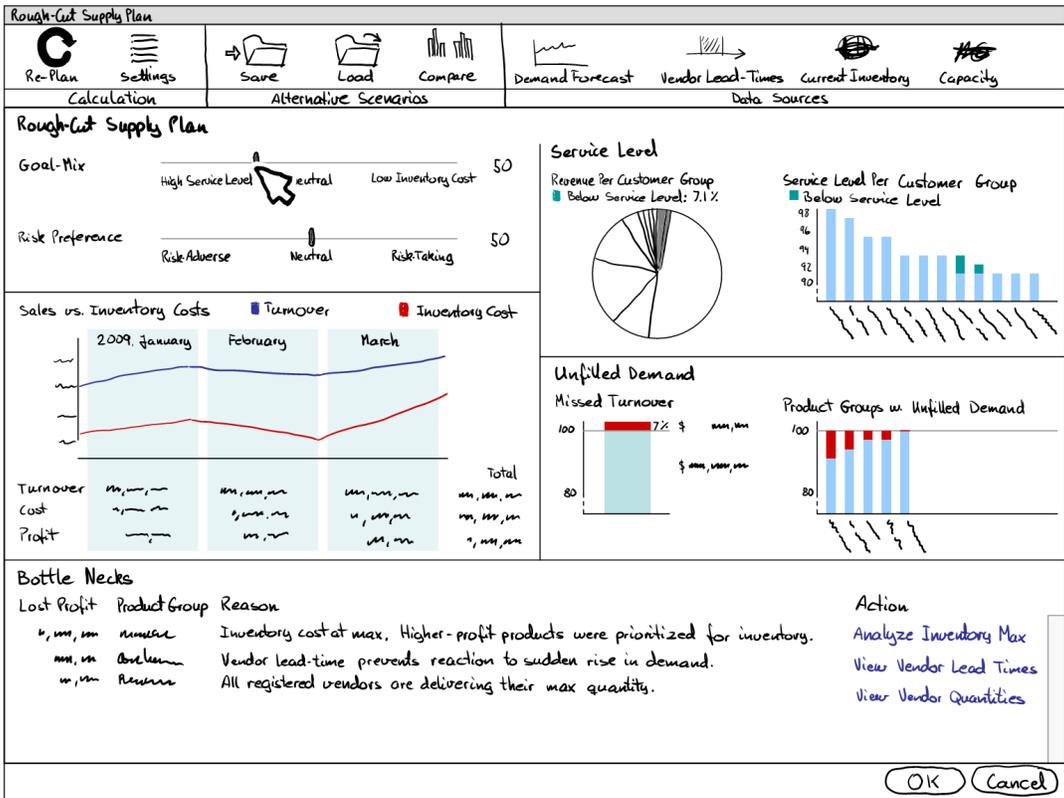
Slide 29



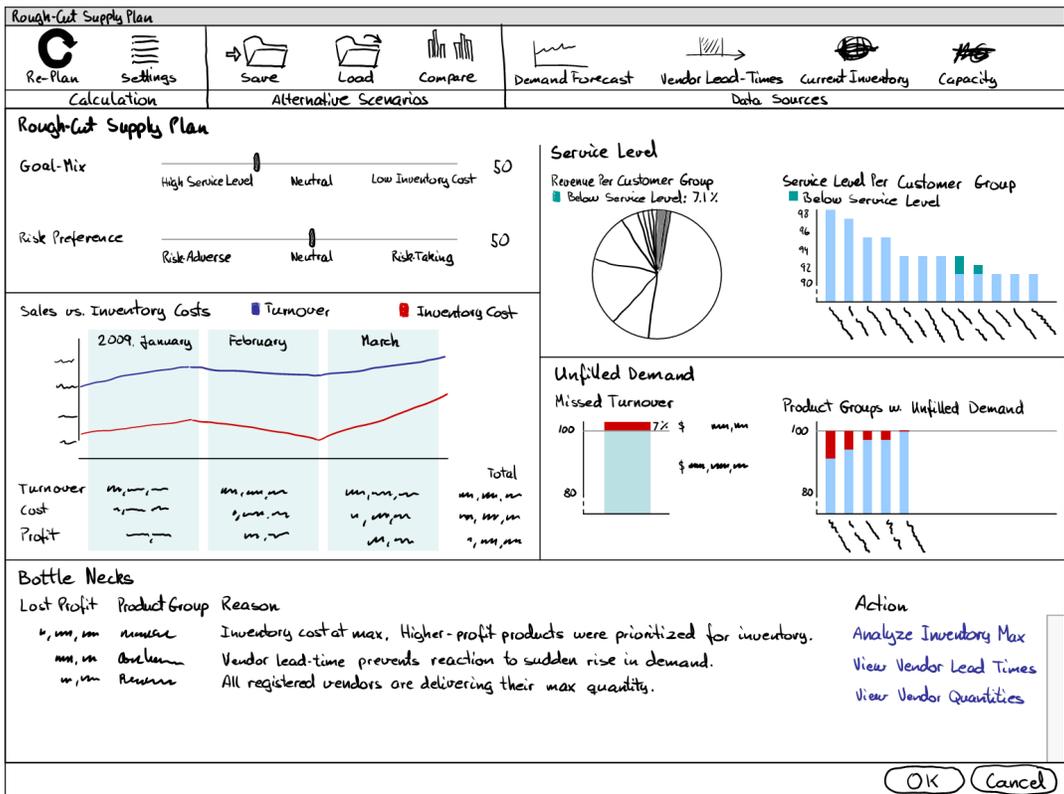
Slide 30



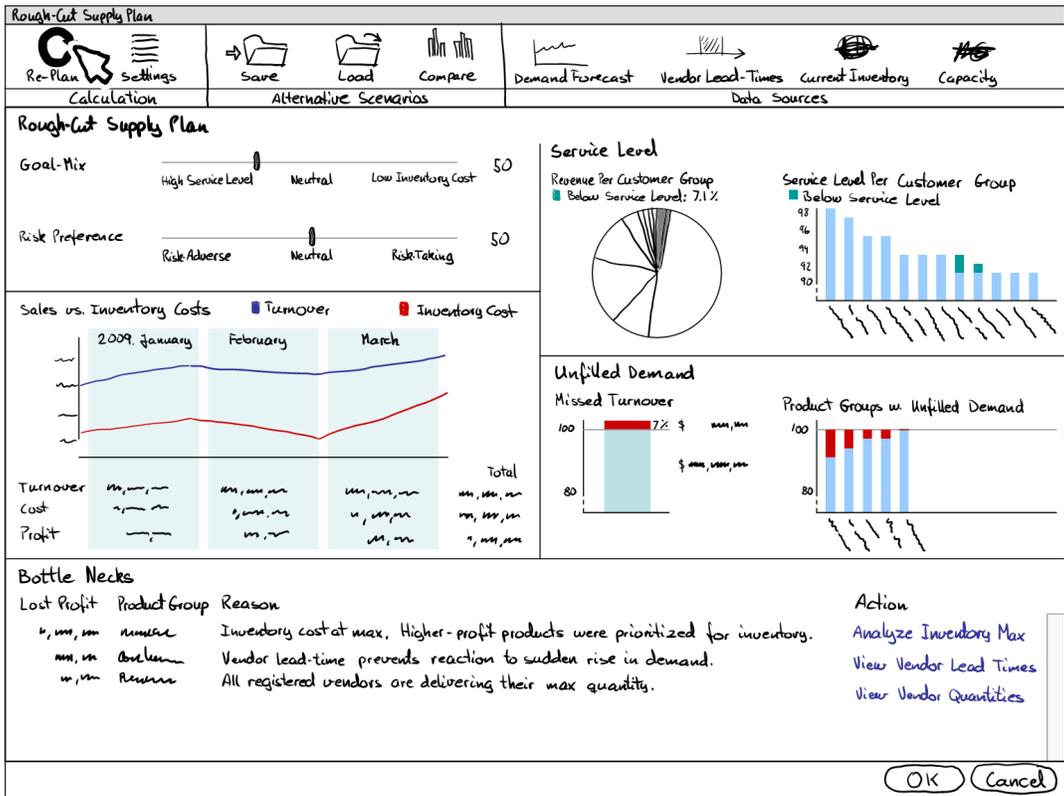
Slide 31



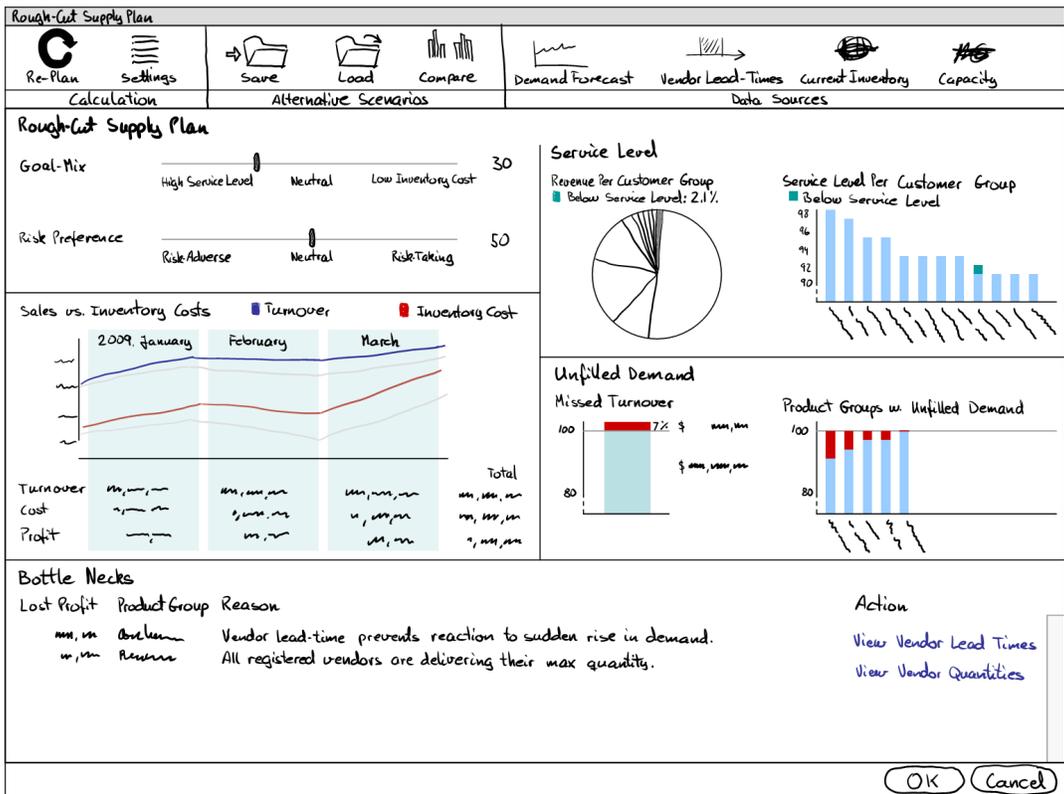
Slide 32



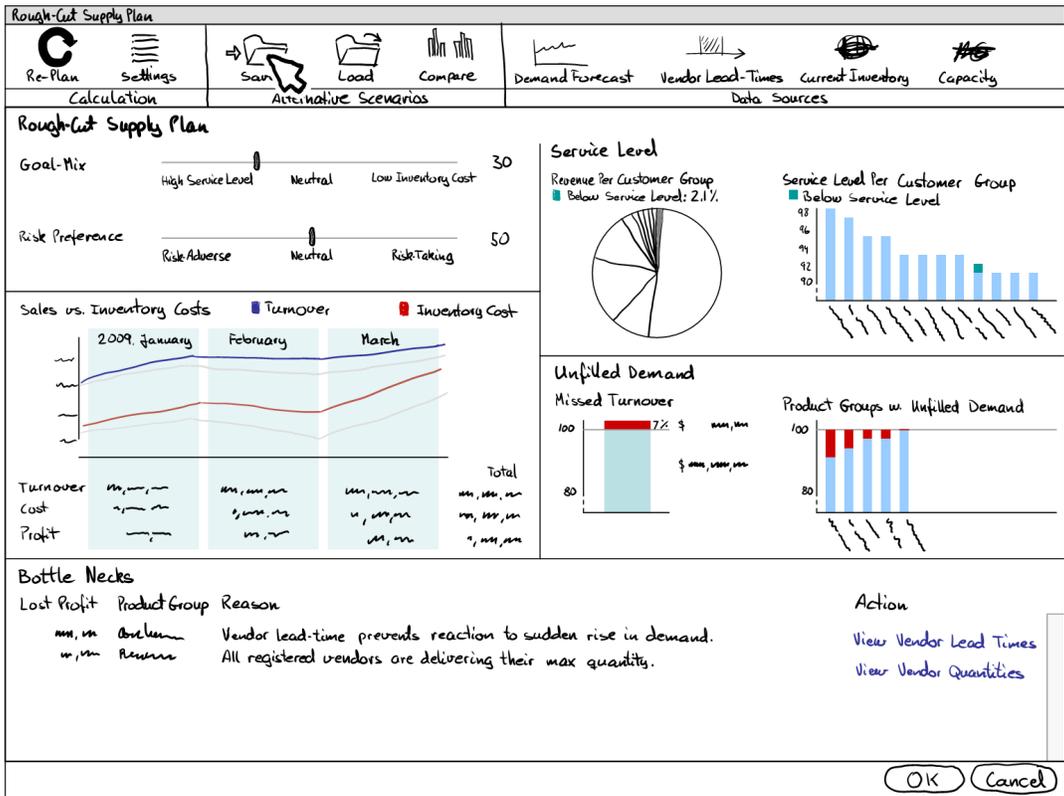
Slide 33



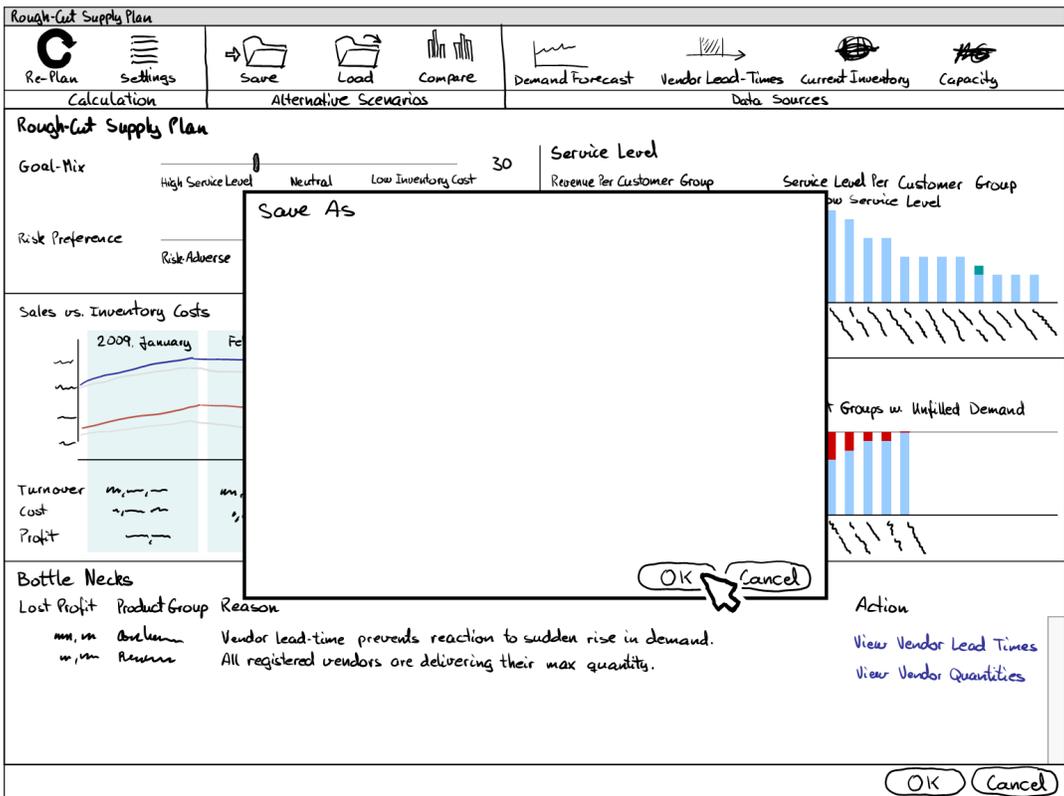
Slide 34



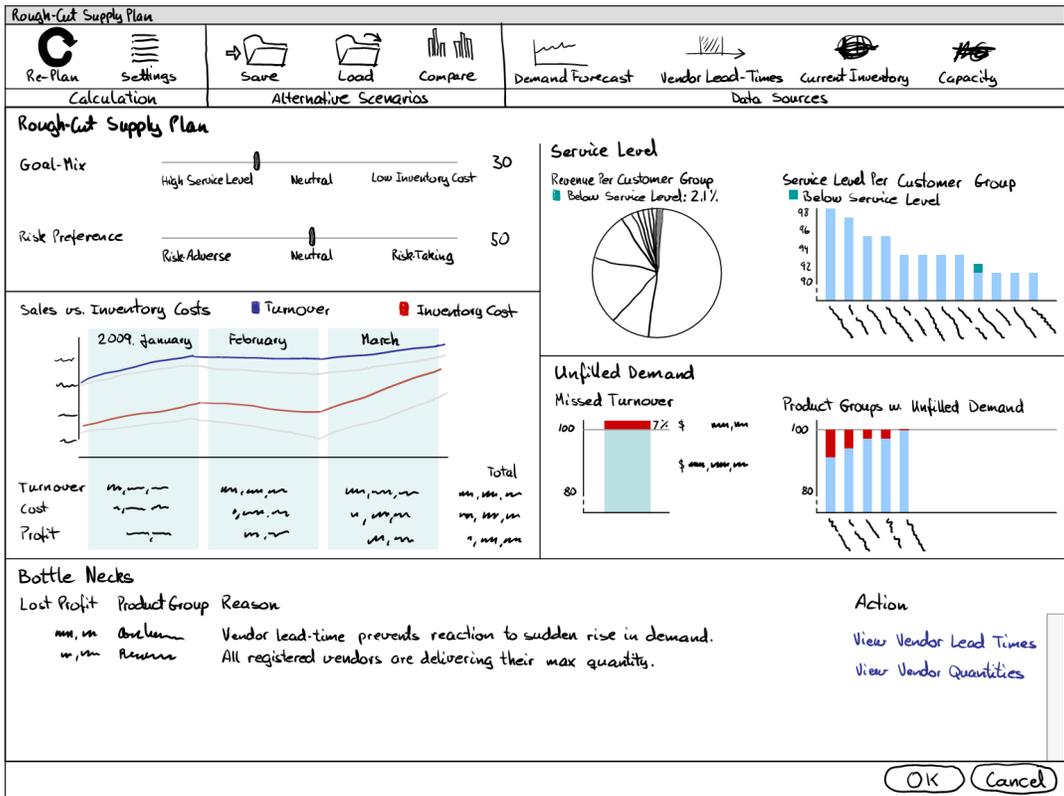
Slide 35



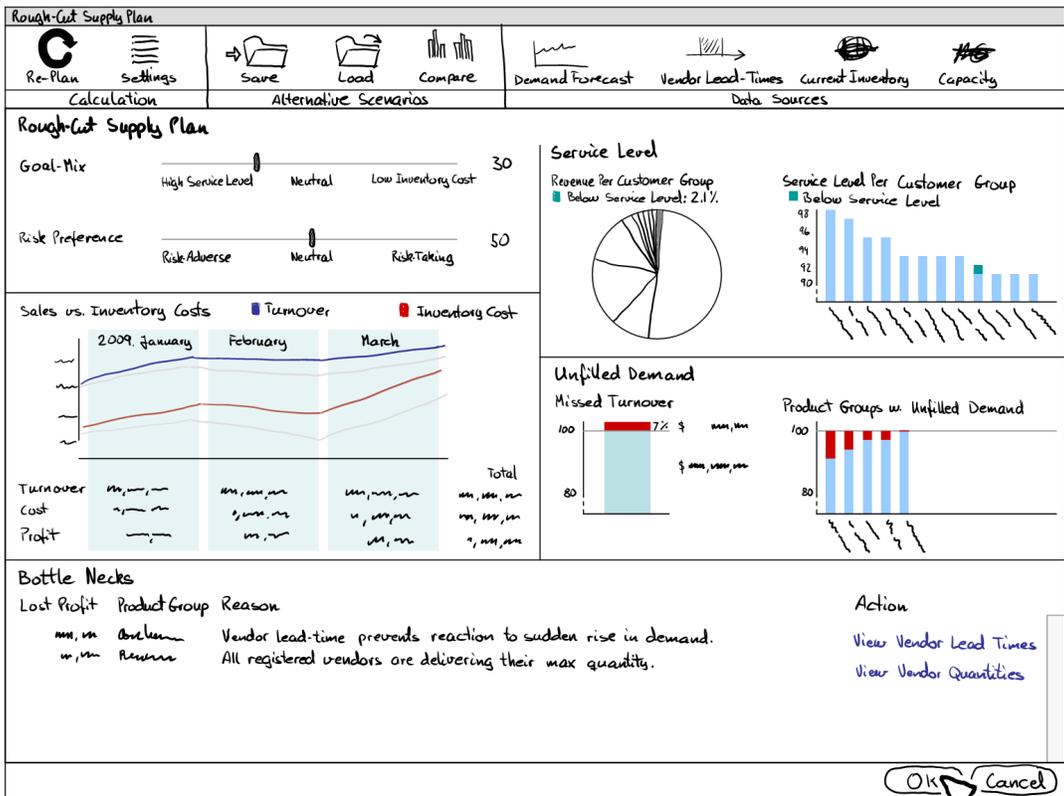
Slide 36



Slide 37



Slide 38



Slide 39