BIGTIME

Predictive Capacity Planning

Using the Past to Predict the Future

Table of Contents

Introduction

The professional services industry is, at its heart, in the business of transformation. Management consultants facilitate the transformation of organizational structures and business strategies, technology consultants enable transformation through digital enablement, and marketing consultants envision the transformation of consumer experiences.

There is one thing that all professional services have in common: their business models are fundamentally about the transformation of skills and knowledge of really smart people into value for their clients — value that clients are willing to pay for. The world of consulting is all about the transformation of talent into money.

What makes professional services and consulting really valuable to clients is that the talent in question is nearly always hard to find, specialized, and expensive. Because of this, successful professional services organizations become successful only if they can figure out that magic formula to transform talent into money in the most efficient way possible. One of the keys to that efficiency is proper capacity planning.

Most services organizations engage in capacity planning in one form or another along a continuum. This may range from a glorified guess to a well-defined resource management workflow to disciplined mathematical models. There's no one-size-fits-all answer for what the right model is for your business. The right approach for you is dependent on factors ranging from team sizes, geographical and demographic diversity, project profiles, and the overall maturity of the firm's processes.

This white paper is intended to help organizations understand that spectrum and decide where along the continuum is most natural for their environment and their specific needs. It also explains a model that we call predictive capacity planning that is targeted at organizations operating in complex resourcing environments or that are managing large teams of consultants. It's for companies that have the basics covered and are looking to take their capacity planning process to the next level.

We hope these concepts around disciplined capacity planning provide valuable perspectives as you think about the transformation of your own internal resourcing and planning processes.

Happy reading,

Brian Saunders

CEO & Founder of
BigTime Software

Steve Chong

Head of Strategy of
BigTime Software

Capacity Planning in Context

Before we dive deep into the mathematics of predictive capacity planning models (and don't you worry, dive deep we will), it's useful to provide some context around what capacity planning is in the first place.

> **Capacity planning** is the delicate art and science of ensuring the right person is available to be put on the right project at the right time.

It's worth noting up front that ultimately, managing a professional services organization is, at its heart, a human endeavor. Manufacturing operations may be all about managing the flow of entirely identical widgets through production lines. Field dispatch environments might optimize the scheduling of completely interchangeable field service technicians. But, as anyone who's managed a professional services team knows, "entirely identical" and "completely interchangeable" are words that rarely describe the professional services environment.

Successfully running your organization involves managing extremely talented and passionate consultants, each interacting with their project teams, clients, and business partners. While the remainder of this white paper leans into the mathematical, logical aspects of resource management, it is by no means intended to diminish the importance of the personal side of personnel management.

The aspects of capacity planning that we'll discuss in this white paper will dive into techniques that can be employed to understand the supply and demand for talent within a services firm:

> **The supply side** is most often modeled as the availability of talent within a pool of resources, whether they are direct employees of the organization or external contractors that are available to be staffed.

> **The demand side** is most often thought of as work that needs to be delivered by the team, whether it is committed work or high-probability prospects.

The underlying goal of capacity planning is to ensure that supply is sufficient enough to meet demand. But, because any hour that is not productively used is immediately gone forever, oversupply is almost as detrimental. While that may sound like a simple goal, achieving it with any level of repeatability, efficiency, and scalability can often be more involved than it appears.

Capacity Planning Continuum

Every professional services organization we've ever encountered engages in capacity planning in one way or another. Unlike other core operational processes that may be more standardized, the level of rigor, discipline, and process maturity around capacity planning probably varies more widely among services businesses.

For smaller teams, typically of a half-dozen billable consultants or fewer, capacity planning is often just handled informally. "Back of the envelope" plans are good enough to give an organization visibility into the future.

That quickly becomes unrealistic as team sizes and complexity grow, even with organizations of just a few dozen resources. Once a services business branches off internationally, increases the diversity of its services offerings or business lines, and grows to thousands or tens of thousands of consultants, it really needs to take its capacity planning capabilities up a notch or two.

Capacity planning at this level requires systems, processes, and people that would be overkill for smaller, simpler businesses. What we often see is an evolution or a continuum that looks something like this:

Pipeline Modeling

The simplest method of capacity planning is probably taking a high-level estimate of committed and high-probability projects in the pipeline compared to an informal assessment of how busy existing resources already are. This often occurs in very small teams characterized by minimal differentiation between types of resources needed, low complexity around the resourcing process, and little need for involved review and approval workflows.

Pipeline Modeling Pros: Good for smaller teams looking for a high-level, back-of-the-envelope look at capacity among similarly skilled staff.

Pipeline Modeling Cons: Minimal checks and balances, may need multiple recalibrations as project scopes and timelines shift.

Backlog Management

Backlog management involves a slightly deeper level of project planning by taking the high-level estimate described above and expanding it into a more detailed, robust plan of who is needed on the project and when. It may entail breaking out the project into individual tasks, assigning people to those tasks, and then aggregating that data into a resourcing plan for the project.

Alternatively, it might remain a high-level estimate of what people and what percentage of those people are needed over the lifetime of the project. The labor estimate is then turned into a budget for the project and work is performed against that budget. What remains is the backlog, or the remaining work committed to the client to complete the project. This takes simple pipeline modeling to the next level of detail for an individual project in order to bring a little more definition to the demand side of the picture.

Backlog Management Pros: Incorporating a budget provides guidelines for project management and effectively accounts for time not yet used.

Backlog Management Cons: Getting a rolled-up idea of organization-wide capacity can be challenging as every project team needs to plan and execute with the same level of detail and precision.

Staffing & Resource Management

Built upon the demand modeled out in the project planning process, staffing and resource management is the process of matching up people and projects — effectively satisfying the demand with supply.

As with other processes, this can run the gamut from simple, informal conversations leading to assignments tracked on whiteboards to fully staffed Resource Management Offices (RMOs) run by experienced Resource Managers complete with formal RMCP certifications (Resource Management Institute, Resource Management Certified Professional, 2022). Most often, this process involves satisfying staffing needs with people based on several inputs, including finding people that have the right mix of skills and availability.

Some organizations think about staffing and resource management on a more micro level, considering each individual staffing decision in isolation. Those that are able to think about this process at a more macro level are starting to work their way into the next level — capacity planning.

Staffing & Resource Management Pros: Puts billable resources at the center of planning, incorporating skills, and balancing availability.

Staffing & Resource Management Cons: Can become a labor-intensive manual process if scalable systems aren't implemented.

Capacity Planning

Capacity planning takes the output of the staffing and resource management process and turns it into a perspective around what types of investment need to be made into the talent base of the organization. That investment might take the form of professional development tracks for existing employees, establishing new commercial partnerships to supplement existing capabilities, or hiring to acquire additional talent in particular areas.

Capacity planning takes the operational and transactional nature of staffing and resource management and transforms it into a more strategic perspective for the business.

Capacity Planning Pros: Getting a view into future demand for skills and staff moves an organization from reactive to proactive.

Capacity Planning Cons: The interdependencies of people, projects, and profits must all be accounted for and data must be reliable and regularly updated for any forward-looking planning to be accurate.

Predictive Capacity Planning

One of the key purposes of this white paper is to introduce the notion of predictive capacity planning, which is a model that takes the concept of traditional capacity planning to the next level. While capacity planning can provide perspectives around the overall supply and demand of resources within the organization, this approach augments it with a predictive analytics model.

Predictive capacity planning effectively allows professional services organizations to predict staffing needs more accurately than by only considering the known demand projected by project planners. It's a technique that effectively uses the past to predict the future.

Predictive Capacity Planning Pros: Data-driven predictions remove subjectivity and more accurately project future capacity and resource needs.

Predictive Capacity Planning Cons: Due to data requirements, it's only practical for enterprise companies with mature PSA software usage and/or data sets

Predictive Capacity Planning Overview

Predictive capacity planning is a model that can be incorporated into BigTime Projector Professional Services Automation software to improve the accuracy of future resourcing needs. Because BigTime Projector tracks the entire history of projected hours for every week by every project manager on every project for every team member, and because it can compare those projections against actual labor usage, it provides for a rich set of data that can be mined to spot patterns.

This model effectively uses data about how accurately and consistently project planners have forecasted their resourcing needs in the past. It then takes that data and converts it into a future-looking probabilistic model that estimates how likely the project planners' forecasts will be to cover actual staffing needs. Finally, it applies this probabilistic model to the project planners' forecasts to determine a predicted level of demand that can be more accurately compared to the supply of talent in the resource pool.

Analyzing Historical Patterns

The process of converting historical data into a predictive model is really the secret sauce that probably requires a bit more explanation. This process starts with looking at deterministic projections of identified resourcing needs modeled in the system. To that, it uses historical patterns to supplement identified needs with a factor that represents how much visibility (or lack of visibility) project planners have had in the past for similar assignments.

Most of the time, this visibility factor is a positive number since actual demand usually exceeds projections. Project planners naturally lack visibility the further into the future they try to project. However, if planners habitually overestimate projected demand, this could, on occasion, be a negative number as well. These historical patterns are also used to add an additional factor that accounts for variability, or how consistently planners have been in their projections of resourcing needs. These all end up summing to a prediction of the actual resourcing needs that is more accurate than just looking at projections of identified resourcing needs:

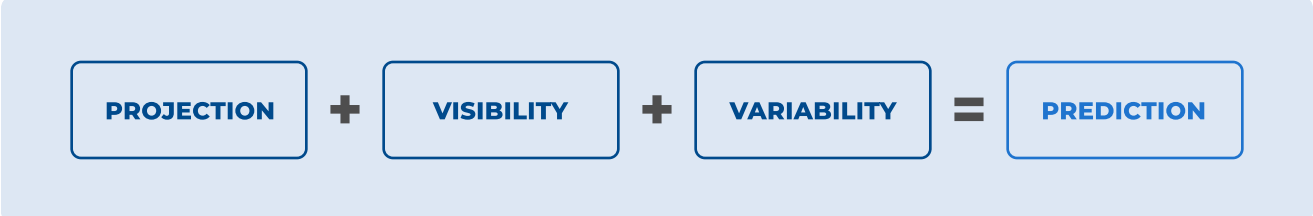

Visualizing the Data

The way this appears visually in a dashboard might look something like this.

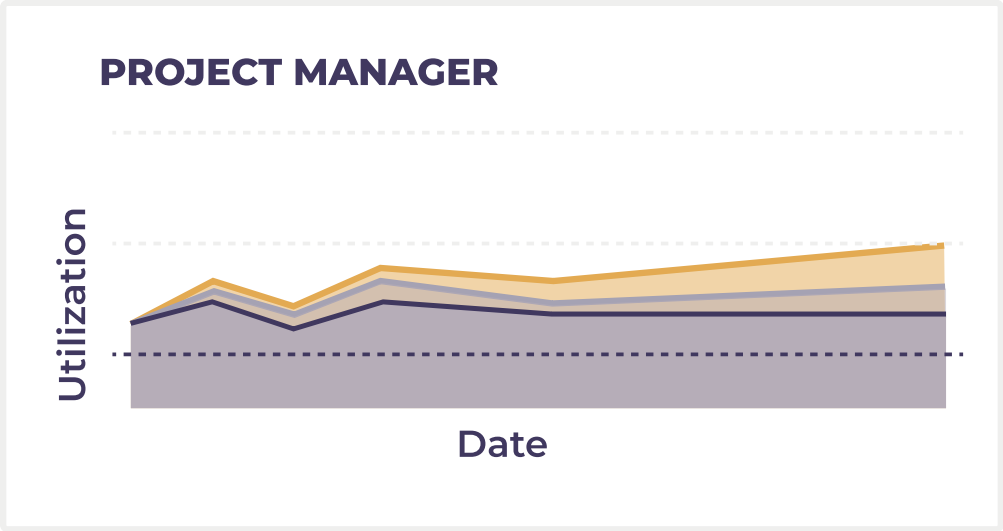

Here, we're looking at a graph of predicted utilization over time for project managers within the organization.

The base starts with the projected resourcing needs. The next layers represent the visibility factor and the variability factor which are based on historical behavior. All of this adds up to the prediction that can be compared to the utilization target (the dotted horizontal line). What's interesting about this perspective is that it gives the organization the ability to understand the difference between true capacity problems and process problems. For the example of our project managers, it's obvious that even the deterministic projections exceed our utilization target, so we have a true capacity problem.

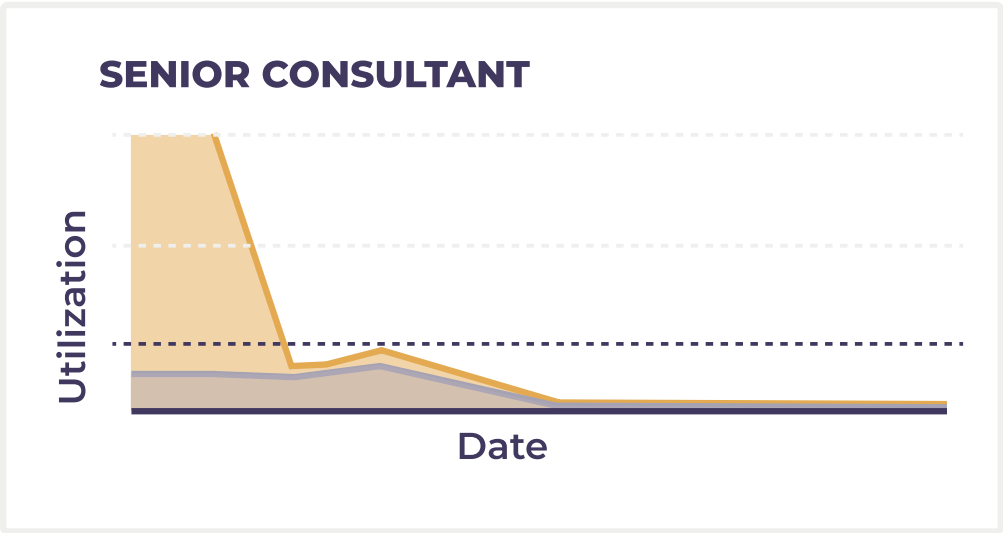

Compare that picture to a similar picture focused on a different type of resource, senior consultants.

Here the identified projected need is far below the targeted utilization level. Based solely on that information, it seems like we have far too much capacity to justify based on the demand. However, looking at the visibility factor, it looks like we're not in as good shape as possible, mostly because project planners have habitually under-forecasted actual demand for senior consultants. What's perhaps more troubling is the large variability factor that indicates that not only have projections been low, but they've also been quite inconsistent. This has resulted in a large "safety staffing" factor that the model is calling for to satisfy the demand.

So, while on initial analysis, it looks like we have too much excess capacity in our senior consultants, in reality, we're much closer to capacity than expected and also need to get our forecasting process under better control.

BigTime Projector Data Models & Dashboards

All of these theories may be interesting in concept but would be useless without practical tools that can be quickly implemented in a typical services organization. To that end, we've built data models and dashboards in BigTime Projector that do just that:

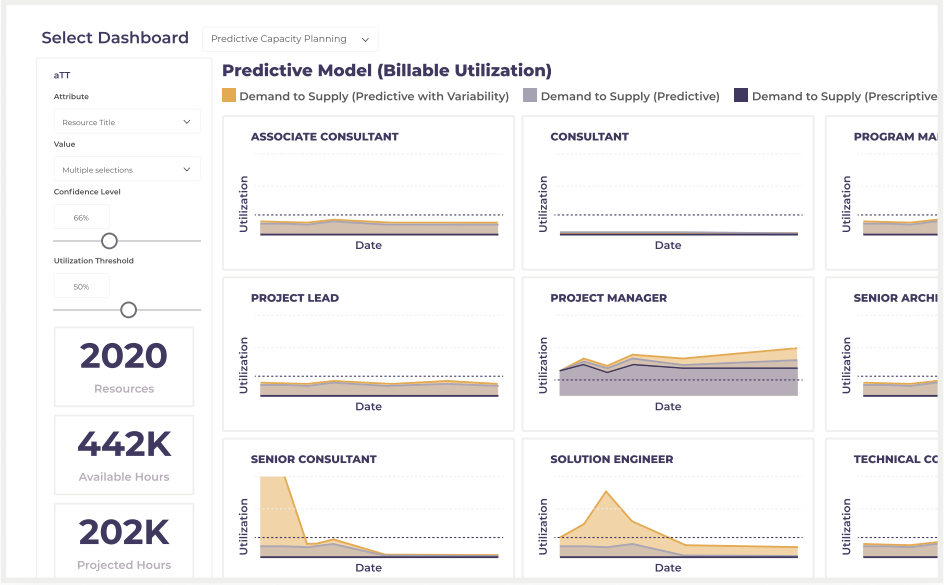

These tools automate the process of analyzing your historical data to spot trends and apply those trends to current projections. They also provide flexibility around how you choose to segment the data as you're searching for those behavioral patterns. In your organization, perhaps there's little difference in the visibility and variability factors from title to title. With these tools, you can quickly pivot to segmenting by geography, practice, business line, manager, or any of the dozens of other attributes that can be tracked in BigTime Projector.

Predictive Capacity Planning Details

While this section will dive a little deeper into the overall approach we've used to build a predictive capacity planning model and have a little bit of the math involved, the point is that organizations don't need to understand all the details behind the model in order to take advantage of it.

So, if you're not interested in all the details, feel free to skip this section. But, for those of you who are gluttons for punishment, we've broken down our predictive capacity planning model into five (relatively) simple steps:

1. Gather Historical Data for Calibration

One of the problems with statistical analyses and machine learning models is that they can be very data-hungry. Not only do they require a lot of data, but it needs to be clean and consistent as well. Fortunately, for the purposes of this predictive capacity planning model, BigTime Projector contains all the underlying data that is required, all laid out in a consistent, well-structured data model.

Gathering this calibration data consists of taking a look at how many hours were reported by people on all of the projects they worked on over the course of a year. The algorithm also looks at the number of hours that were forecasted for each of those people/project combinations (called roles in BigTime Projector) a week before, two weeks before, three weeks before, etc. Finally, it gathers attributes that describe each of those roles, ranging from who was the project manager forecasting the resource needs, what sort of person was required, what part of the organization they came from, and so on.

2. Perform Statistical Analysis of Calibration Data

In a decent-sized services organization, this calibration data may range from hundreds of thousands to millions of data points, which should be sufficient to do some reasonable analysis on, even if we need to do a few rounds of segmentation.

This analysis consists of generating two main perspectives about the relationship between the forecasted hours and the actual hours worked that are loosely based on some concepts from the world of quantitative stock analysis.

The first perspective is based on the idea of a stock's alpha (ɲ) which describes whether the stock performs generally better or worse than the market as a whole (Investopedia, Alpha, May 2021). In the context of our predictive capacity planning model, our ɲ is used to describe whether the projected demand was consistently over or under actual resourcing needs. This is roughly equivalent to the visibility factor we discussed earlier. Unlike a stock's ɲ, which is a single parameter, our ɲ is a function that changes over time—that is, ɲ(t). This makes sense when you think about it as project planners generally have pretty good visibility of what resourcing needs are next week, that is ɲ(1), as compared to needs 10 weeks in the future, or ɲ(10).

In quantitative stock analysis, an investment's beta (ɴ) is a parameter that describes the stock's volatility as compared with the market as a whole (Investopedia, Beta, Jan 2021). Our predictive capacity planning model uses ɴ to describe how consistently project planners forecasted their resourcing needs — basically, the level of variability mentioned earlier. Even if, on average two project managers forecasted their staffing needs exactly but one sometimes significantly under-forecasted and sometimes over-forecasted, they both would have similar ɲs but different ɴs. And, like the ɲ(t) function, ɴ differs from the single stock market parameter by having a function ɴ(t) that changes over time in our predictive capacity planning model.

3. Transform Statistical Analysis into Probabilistic Model

Once we have the statistical analysis taken care of and have calculated our ɲ(t) and ɴ(t) functions, we can turn that around and transform that into a probabilistic model. This starts with asking the organization for a confidence level that they want to cover. That is, since project planners are not at all perfect in projecting their staffing needs, given a particular level of variability — what percentage of time is necessary to be able to satisfy the demand and how frequently is it acceptable that the organization can't cover it?

This is a concept that loosely hails from the worlds of supply chain and operations management known as safety stock. If you are able to measure your average rate that a particular product is purchased as well as how consistent (or inconsistent) that rate is, you can calculate how much extra you may need to keep in inventory to make sure you don't run out x% of the time.

Of course, that leads to the obvious question of what is an appropriate x%? In general, we recommend that most organizations keep this relatively low in the 60%-70% (one sigma) range rather than the higher certainty 95%-99% (two to three sigma) ranges. This will vary from situation to situation and really depends on the implications of not being able to cover demand in a particular area weighed against the cost of having sufficient "safety staffing" to ensure no role goes unstaffed. A task may need to slide to the following week or a consultant may need to put in some overtime, so a fairly loose confidence level is typically reasonable.

With this confidence level in mind, we can use some transforms common in probability and statistics to calculate what's known as a Z score. This Z score can then be used in conjunction with the μ(t) and σ(t) functions to calculate how much of a "safety staffing" level needs to be maintained.

4. Apply the Probabilistic Model to Create a Predictive Model

With this concept of the Z score in place, we can finally start applying it to the organization's real forecasts to start building something that is more familiar and more usable for business users. The Z score can be used in conjunction with the μ(t) and σ(t) functions to figure out how much the organization's deterministic resourcing projections need to be modified to cover predicted demand.

For instance, say our project planners have projected needing 10,000 hours of work to be done by technical consultants three weeks from now. We know that on average, our μ(t) function indicates projections three weeks out tend to be 20% low. With this information, we can project we might need enough resources to cover 12,000 hours. In addition, based on historical behavior, that 12,000 hours may actually vary between 9,000 and 15,000. In order to cover demand 66% of the time, our σ(t) function and Z score indicate that our best prediction of demand may end up being 13,000 hours.

5. Incorporate the Predictive Model into Practical Tools

Finally, with the predictive model in place, we need to wrap all the concepts, statistics, and probability into practical tools that can run these calculations quickly, intuitively, and at scale.

We've done exactly that in our predictive capacity planning model that can be incorporated into BigTime Projector. With it, users can transform their deterministic projections of resourcing needs into more accurate predictions. These tools also give users the flexibility to explore what the effects of increasing confidence intervals are, whether predictions are truly affected by insufficient resources or by out-of-control processes, and which teams/locations/types of resources need the most attention.

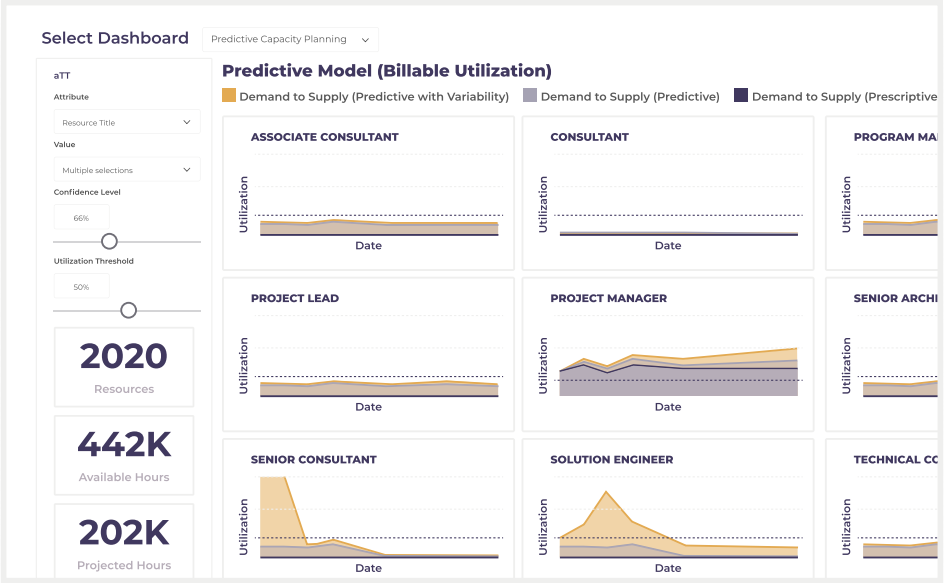

This model can also be extended into a custom machine-learning model by incorporating a periodic refresh of the underlying calibration data. By pulling in new historical data on a rolling basis, the system will continually tune itself over time to reflect changes in process, teams, or business models. All of this together can transform what may have started with glorified guesswork into a mathematically rigorous model that adapts as you improve.

Conclusion

As you can see, the idea of accurate capacity planning is a critical one for talent-hungry professional services organizations. Depending on the organization's needs, that sort of planning can range in a wide gamut from very informal to very disciplined. It can also vary in nature from the largely qualitative to rigorously mathematical. The point is, whether you're planning on time horizons of 5 days, 5 weeks, 5 months, or 5 years, there's a useful capacity planning model that you should start with and continue to improve upon over time.

For those organizations that are running larger and more complex teams, are looking to plan more strategically over longer time horizons, and already have the basics covered, moving from a deterministic projections model to a predictive capacity planning model can provide significant benefits.

1. Leverages historical behaviors to predict the future
2. Incorporates visibility and variability to improve accuracy
3. Differentiates true capacity constraints from process problems
4. Adapts over time to reflect process improvements

In order to implement an effective predictive capacity planning model, organizations need to have access to data — data that is accurate, well-structured, and plentiful. They also need to understand how to leverage that information to transform it from raw data into insight and historical trends. These trends then need to be applied to future resourcing projections in a model that is mathematically rigorous and contextually relevant. Finally, this model needs to be incorporated into practical tools that businesses can leverage without having to understand the details of statistical analysis and probabilistic modeling.

We hope you found this explanation of capacity planning concepts useful and wish you the best of luck with your efforts to continually improve your own capacity planning capabilities. As always, if you're a consulting or services firm looking to optimize every aspect of your operations the team at BigTime is here to help.