WHAT IS A CONTROL TOWER?

AN OVERVIEW OF THE FUNDAMENTAL CONCEPTS AND PRINCIPLES BEHIND ENTERPRISE CONTROL TOWERS

January, 2020

Kelly J. Thomas

worldlocity

Introduction

"Control tower" is a term now widely used in supply chain management to describe a variety of capabilities across a variety of domains. The term is so widely used that there is an implicit assumption that others know what it means. In fact, those using the term may not have a solid definition, or may have a definition that is completely different from those hearing the term. This leads to a sort of false consensus, which can cause inflated expectations and inconsistent understanding within and across organizations. This happens a lot in the world of software and technology and is part of the natural evolution described by what Gartner calls a "hype cycle."

Adding fuel to the control tower discussion is the digital revolution, which is really just the next chapter of an evolution that has been unfolding for fifty years. In today's business world, you can make anything sound fancier by adding the prefix "digital," as in "digital control tower." This is technology's equivalent of ketchup's "new and improved." As in the case of ketchup, there is some reality associated with using the term, and there is also a healthy dose of marketing.

There are certain foundational principles that are common to all control towers. It is useful to explore and define these to accelerate understanding and to flatten the hype cycle. This article discusses the importance of control towers and their common foundational principles.

High in the Sky

Analogies and metaphors are useful in learning and in taking ideas from one discipline and applying them to other disciplines. Humans are really good at this. Applying old, well-understood patterns and words to new technologies and disciplines helps accelerate thinking in the new area.

"Control tower" is a term borrowed largely from the airline industry. In describing its application to supply chain problems, the most commonly used icon in PowerPoint and other documents is an air traffic control tower standing high in the sky, a combination of people and technology, managing the time-phased inflow and outflow of planes, passengers, and cargo.

If you apply the air traffic control tower functional analogy to supply chains, you find that these capabilities have been present for years across various disciplines, including manufacturing, distribution, and transportation. For example, a manufacturing plant has a view of its incoming material, its work in process, its finished goods, its machine and labor status, and its orders. It has a master production schedule to profitably satisfy demand. Plant management orchestrates the plant in order to satisfy demand. Similar patterns hold true for distribution and transportation. This orchestration capability is essentially a control tower. So, control tower capabilities – if not in name – have been around for a long time.

Transportation – movement and flow of goods from one place to the other via ground, air, and sea – is the area of supply chain management that most closely correlates to the idea of a control tower in the air traffic controller sense. That is to say, it involves transportation assets that move things. Thus, transportation was first to adopt control tower terminology in its operations, a couple of decades ago.

Mission Control, Control Rooms, War Rooms, Command Centers, Control Towers

Control tower terminology is also used interchangeably with other terms, including control room, war room, command center, and mission control. At their core, all of these constructs are similar. **Their fundamental objective is to control a process to effect an outcome**. In order to do this, they need visibility into what's happening in the process, along with resources and capabilities that can directly intervene in the process to keep it operating along its intended path.

War rooms, or their equivalent, have been used for centuries to direct armies in war. In business, they have historically been set up to deal with crises; they are set up to micro-manage a crisis (or other business challenge) and remain in place until the crisis abates and the situation is back under control. Production control rooms have been used in manufacturing facilities at least since the 1920s.

In the past twenty years, a significant trend in control constructs is that they have steadily moved from focused operations or facilities into having a much larger span of control. Thus, control towers (or control rooms) have evolved up and out, as shown in the below figure. For example, they have evolved up from managing a single facility in a single functional area (for example, a single manufacturing plant) to all manufacturing plants. Likewise, they have evolved out from a single functional area like manufacturing to include suppliers, distributors, and retailers. The net result of this is that control towers are increasingly managing the end-to-end supply chain of manufacturing, distribution, retail, and customer service.

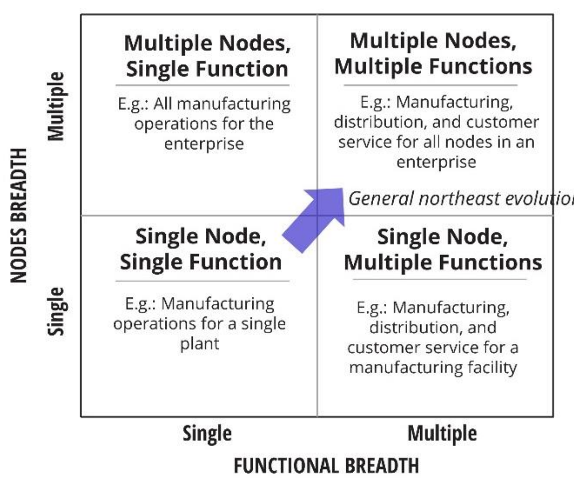

There is also a third dimension that needs to be added to the simple 2x2 shown here: the time dimension. The time dimension considers whether the control tower focuses on operational (execution), tactical, or strategic decision making. This will be discussed later.

Why the Hype?

If control towers and their like have been around in some form for a long time, then why is there so much hype now?

About twenty years ago, Gartner introduced its hype cycle to describe the evolution of new technologies, particularly information technologies. If you look at Gartner hype cycles from ten years ago, you find that a good percentage of the hype was exactly that – hype. In other words, the things included in the hype cycle did not live up to their initial hype in the ensuing ten years; some fizzled out, some morphed into something else, and some were simply evolutionary. On the other hand, a good number of them developed into useful, widely-used, value-driving capabilities.

And, hype does serve a purpose in the world of technology – it fuels investment. Ideas get hyped and then investment flows to those ideas. A certain percentage of the hyped items actually produce business value; the others still serve a purpose because investment drives discovery that in many cases leads to other ideas, concepts, and technologies that do deliver value (and in turn, creating new hype, driving a self-reinforcing cycle).

While there is a fair amount of hype surrounding control towers, this hype is fueled by real and important business needs. **Control towers have become increasingly important in the past ten years because the competitive environment places increasing value on precision.** What do I mean by precision?

Precision is characterized by increasing granularity in time, product, and service. This is driven by customers increasingly demanding a product tailored to their needs, delivered to them at a time and place tailored to their needs, at a price tailored to their needs. Thus, precision in business means increasingly tight tolerances placed on matching products and services to individual customers. This manifests itself in the supply chain with increasingly tight tolerances on manufacturing, distribution, and retail operations. For example, ecommerce delivery times have shrunk from a week to days to two-days to single-day to same-day. Likewise, upstream supplier on-time in-full (OTIF) windows are shrinking from a week to several days to a single day for greater than 90% of orders over a monthly period.

Increasing precision, as defined here, leads to increasing complexity, which ironically is the enemy of precision. For example, SKU counts and product variants need to increase in order to provide more precise customer choice. This increased product complexity must be managed across the same asset base in order to maintain profitability and return on investment (see Consumer Choice, Productivity, and Managing Supply Chains). For example, consumer goods companies can no longer depend on growth from a few well-established mass-market products; they must seek growth from niche products tailored to the needs of small segments. The beer market is a good example of this:

there are now more than 7000 breweries in the U.S., versus fewer than 80 in 1978.

Complexity will continue to increase, driven by a voracious human appetite for better, faster, cheaper. This appetite fuels an ever-faster-moving, self-reinforcing process in which companies introduce nice-to-have features and capabilities that quickly become must-have features and capabilities. The time gap between when nice-to-have features are introduced and when they become must-have features is shrinking. For example, once two-day delivery was introduced, it quickly became an expectation; likewise, when car backup cameras first became available they were a luxury item but within several years they became a standard feature.

Staying ahead of this curve requires continued investment. Over the course of the past ten years, the business trash heap has been filled with companies that either could not invest, didn't want to invest, didn't invest fast enough, or invested in the wrong things at the wrong times.

Control towers are a means by which to provide precision in an increasingly complex business environment. This is the key reason they have become a hot topic in the past ten years.

Delivering precision also means dramatically lowered time cycles for sensing and responding. A control tower that spans the entire enterprise provides an aggregate up-to-the second view of operations. In the past, such aggregate views had to be assembled in after-the-fact reports. The lag time between an aggregate view and current time was followed by an equal lag in response, creating a tail chasing dynamic characteristic of the bullwhip and similar effects. A control tower couples real-time visibility with real-time response. (For a formal definition of "real-time," see "What is Real-Time?").

Boundary Conditions

Precision is typically lost at organizational boundaries. And, this precision loss has a sequential compounding effect, as information and materials flow across organizational boundaries. **This leads to a fundamental law of supply chain dynamics – the more organizational boundaries that are involved in the fulfillment of an order, the higher the chance of loss of precision**. This compounding effect was studied in detail in Jay Forrester's 1961 work "Industrial Dynamics," and is a key contributor to the commonly-known bullwhip effect. This is where modern control towers come in.

In the airport control tower analogy, it is clear that an outage or problem at a single airport can have a cascading system-wide impact. For example, the impact of a winter storm in Chicago can be felt by travelers in New York, Atlanta, and Los Angeles. In the business of managing airplanes and passengers, there are not many degrees of freedom to handle such problems. Having said that, if there is some ample advanced visibility, then assets can be positioned to minimize the impact. The same is true in supply chains.

Supply chains are complex, multi-organizational systems, with each organization having independent objectives. **Attaining a precise output of such a complex system cannot be attained by simply attaining precise outputs from its components and then stringing the**

components together. Thus, control towers in recent years have increasingly focused on coordinating operations across functional boundaries. For a manufacturer, for example, this might mean inbound materials, conversion operations, and outbound distribution.

What Are We Trying to Control?

Control towers are based on a closed-loop control paradigm, in which the tower senses what is happening in the real world, identifies deviations (or gaps) from an objective function, and then responds to eliminate the deviations (see the diagram below). In the process, it looks at various alternatives and chooses the one that best closes the gap, and then places that alternative into action. In order to evaluate alternatives, the control tower may run simulations or scenarios and then add human judgment to identify the best changes to make to eliminate or mitigate the deviation. A scenario is constructed by taking the current state of operations and changing various inputs; the scenario is then run to understand how the changed input variables affect output variables. In today's world, perhaps the most important control tower capability is the ability to rapidly execute and consider scenarios.

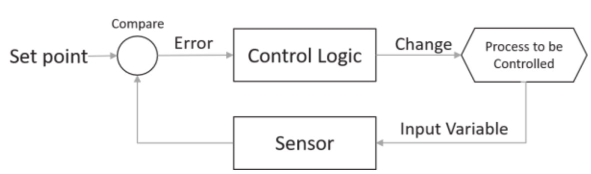

At a macro-level, the objective function of the enterprise is represented by its plan for revenue, growth, profit, cash flow, and return on capital. Thus, the control tower attempts to achieve this objective function by orchestrating and controlling operational variables within upper and lower control limits or bands. I explored these concepts in a 2009 Supply Chain Quarterly article titled "A Rudder for Course Correction."

Unlike process control, in which a single controller may be looking at a few variables, and only in its current state, supply chain control towers must look at numerous interconnected variables across a time continuum, from seconds to minutes to hours to days to weeks. For example, let's say that orders for a key customer are projected to be late a week from now. There are a set of paths by which a company has solved this problem in the past. This set of paths, or rules, can be codified into "decision playbooks." The late order can be resolved by a range of methods, including shifting inventory from one location to another, shorting another customer, expediting a shipment, or adding overtime to a factory or warehouse. Each of these resolution methods will result in different outcomes for the overall plan. For example, running overtime will solve the immediate problem, but may negatively impact the financial plan for the month or quarter.

Much of the daily activity in supply chains is about managing gaps: inventory is too high or too low, sales are too low, production throughput is too low, warehouse throughput is too low, supplier shipments are late, customer orders are late. In all of these cases, supply chain professionals use intelligence to try to close the gaps. Ideally, this intelligence includes advanced visibility so that the gap can be closed before it realizes itself in the real world and threatens the revenue and profit plan. Thus, much of what is done in supply chain management essentially incorporates closed-loop control. This is the essence of a control tower.

When Did you Know?

When operational problems arise in companies, a common question to ask is “when did you know?” Behind this question is a foundational principle of control towers, indeed any control loop. When and what you knew is the sensing part of the control loop.

Let’s say sales for a certain product line are below plan. It is very likely that someone knew fairly immediately that there was a problem, while high level managers may have only recognized the problem two weeks later. This occurs when information and decisions cascade from one level in an organization to another.

On black Friday in 2005, Mike Duke of Walmart was quoted as saying “by 8AM we knew we had a problem.” This meant that by 8AM on black Friday, management could see that the pace of sales was such that the plan for the entire holiday season – and by extension the yearly plan – was at risk. Given this information, Walmart spent the next five days putting in place remediation actions to get back on course. These course-correction actions resulted in Walmart not only meeting their holiday plan, but exceeding it. This is essentially what a control tower does, through a combination of people, process, and technology.

Keep in mind that this story is from fourteen years ago. The technology available today, along with the operational precision required, have advanced dramatically since then. At that time, this was a sort of war room capability put in place to micro-manage the critical holiday season. **Today’s competitive environment demands operational precision that can only be attained by employing such capabilities every day across an enterprise.**

Much of the discussion today centers on AI-based decision making. Can decisions like these be put on automatic pilot, as advocates of the “self-driving supply chain” claim? The answer lies somewhere between yes and no. It is more likely that in the short-term AI will be used for more mundane decisions focused on what some have termed “taking the robot out of the human.” It will also be used to automate various inputs into the decision-making process.

An AI can tell you that you have a problem such as in the Walmart case. But then again, you probably don't need an AI for that. What is different now is that the AI might be able to tell a week before black Friday that you are going to have a problem (instead of at 8AM on the day of black Friday). It might also be able to help you sort through different permutations of response. However, after all the data is sorted, human judgment must be applied to make a decision. The reason for this is that higher level decisions in supply chains are not binary decisions between one path or another, each with single-variable outcomes. Most higher-level decisions are between multiple paths, each with different multiple-variable outcomes and different risk profiles. And, more importantly, there are probabilities involved, meaning there is a range of possible outcomes within each path.

Presumably, you could program the algorithm with your appetite for risk, but it is unlikely for the foreseeable future that the algorithm would be able to understand things not able to be put into the model, including the geopolitical environment, up-to-the-minute competitive moves, and changing consumer sentiment. In this case, it is better for the algorithm to provide its assessment and for humans to apply judgment to make the final call.

What’s Changed in the Past Ten Years

There are three main components to the closed loop control diagram shown earlier:

1. Process to be controlled
2. Sensing function
3. Response function

All three of these have changed significantly in the past ten years, spurring a growing importance in control towers and similar constructs. Each will be examined here.

A short list of omnichannel fulfillment possibilities

1. Buy from local store, traditional in-store shopper transaction
2. Buy online, ship from local store to home or other location
3. Buy online, ship from remote store to home or other location
4. Buy online, pick up in store (BOPUS) from local store through existing inventory
5. BOPUS from local store through inventory shipped from DC
6. BOPUS from local store through inventory shipped from another store
7. Buy online, deliver from closest DC to home or other location
8. Buy online, deliver from further DC to home or other location
9. Buy online, deliver from DC, based on ATP from factory delivery

First, the process to be controlled has become more complex and interconnected, driven, as discussed earlier, by increasing customer demands for precision. While precision has historically been lost at organizational boundaries, the past ten years has seen significant growth in the need to synchronize organizations in order to fulfill orders. Take omni-channel fulfillment as an example (see “10 steps on the path to omni-channel profitability”). There are now myriad options for order fulfillment; these options involve the commingling of physical and digital channels, as shown in the sidebar at the right. The net result is that there are now multiple virtual supply chains that crisscross the same physical asset base. This requires dynamic node-to-node relationships that must operate with efficiency at scale. Control towers play a critical role as the orchestrator of dynamic virtual supply chains; they help choose the best dynamic path to satisfy orders and then keep the orders on track using closed-loop control thinking described earlier.

Secondly, the ability to sense what is happening in the process has improved dramatically. This improved ability has been driven by IOT direct connections, in-memory databases (digital twins), visualization technologies (including mobile), and event-based data pipelines. For example, a real-time location signal from a truck is a common capability today. This signal can be fed directly to an in-memory database and visualized immediately to a control tower operator. All of this has been enabled by the rapid growth of the cloud and open source projects. Furthermore, digital twins now more closely resemble the physical world through the use of machine learning algorithms that analyze large amounts of data and provide statistical probabilities for data elements such as lead times and yields; previously these data elements were viewed as static, leading to model drift.

Thirdly, the ability to respond intelligently has improved significantly. This is due to improvements in standardized playbooks, scenarios, probability analysis, and decision algorithms. For example, control towers now have the ability to rapidly run many scenarios across a range of probabilities in order to choose the best response path. Additionally, these technologies are far more usable today through dramatically improved usability, visualization, and collaboration. Response speed has also improved significantly, driven by in-memory computing, running at scale in the cloud.

A few words about IOT

A precursor to IOT was RFID, which received a lot of fanfare, but didn't really take hold, largely because of cost. The widespread availability of low-cost sensors, driven by Moore's law and economies of scale, reignited the movement, which is now known as IOT. This time, it is taking hold, albeit at a much slower pace than that predicted by its advocates and by many industry analysts. The pace at which it takes hold for industrial settings is gated by the installed base of industrial equipment. A very high percentage (most reports indicate 70-80%) of IOT implementations involves retrofitting existing equipment, versus totally new equipment, which comes already outfitted with internet connectivity. This is because most industrial equipment has a long lifespan, typically in the 8 to 25 years range, along with a sizable installed base beyond that range. When this generation of equipment is completely flushed from the system, IOT will become ubiquitous, since it will be installed natively on new equipment. The same is true of many household devices which have long lifespans, including refrigerators, washing machines, dishwashers, water heaters and furnaces.

For supply chains, the area of fastest IOT growth is location sensing. Because it is relatively easy to retrofit a truck with a location sensing device (including the use of driver smart phones), location sensing capability (at least at the transportation equipment level) has become ubiquitous. This ability to understand, with a high level of precision, where shipments are, has helped improve order fulfillment precision in the past ten years.

Decision Horizons

Control towers, as discussed earlier, have evolved from single functions and single facilities to multiple functions and multiple facilities, in some cases helping manage the entire end-to-end supply chain. They have also evolved from a mostly operational level focus to higher level decisions and analyses. This brings us to a discussion of "decision horizons."

Function / Horizon	Supply	Production	Warehouse	Transport	Orders
Strategic (Years to months)	Where to source	Where to produce	Where to distribute	Where to transport	Where/how fast to fulfill
Tactical (Months to days)	How to source	How to produce	How to distribute	How to transport	How to fulfill
Operational (Days to milliseconds)	Manage flow	Manage production	Manage distribution	Manage transportation	Manage orders

Historically, enterprise decisions have been categorized into three areas: strategic, tactical, and operational. These "decision horizons" correlate to the potential impact that their respective decisions have on the enterprise. For example, operation-level decisions are fairly confined in that their impact is limited to a focused area of the supply chain, like a warehouse, or even an area of the warehouse like a receiving or loading dock.

Strategic decisions have the ability to impact the enterprise for multiple years into the future. For example, an acquisition of a company to enter an adjacent market can fuel additional growth; conversely, if it is not done correctly, it can distract management from its core business and lead to a multi-year decline. Likewise, the decision to fund a new product program, if successful, can fuel growth for years; on the other hand, if the product misses the market, it can lead to decline.

In supply chains, strategic decisions are about assets: where to place assets, the parameters within which assets should behave, and how assets should interact with each other. Tactical decisions in supply chains are about how to operate assets in order to achieve weekly, monthly, and quarterly plans. Operational decisions are about managing assets in order to achieve daily plans.

Of course, decision horizons are not mutually exclusive: operational decisions may have an impact on tactical plans, and even company-level strategy. These problems, if they are not dealt with for whatever reason, including poor decisions, can result in larger problems that ripple across the supply chain. And, because of lead times, such operational problems only manifest themselves as larger problems weeks or even months later. Tracing the bigger problem back to its root cause can be challenging. Nowhere is this problem felt more than in product recall problems, such as in the case of a tainted food product. It may take weeks or months to discover such execution level problems and even then the source may be not be completely certain.

Thus, decision horizons increasingly represent a seamless flow through time. Furthermore, the frequency of analysis at each level of the decision stack has been increasing for the past couple of decades. For example, strategic decisions can no longer just be made and then put on auto pilot. Consider today's geopolitical environment with increasing trade barriers and volatility on a weekly basis. Those companies that have continuous strategic decision processes (versus event-based decision processes) are at a competitive advantage, since the decision apparatus is already set up and has many cycles of learning (including codified decision playbooks).

Form and Function

So, what does a control tower look like? What form does it take?

From a technology perspective, the control tower is a physical thing. It is made up of hardware and software that delivers information to inform decisions. But a control tower is much more than just technology. It is a set of capabilities, delivered by people, processes, and technologies to assist in controlling operations to achieve a plan. (The plan can be developed by the control tower or can be an input to the control tower).

There are typically multiple control towers within an enterprise. The operating committee, the S&OP process, and the transportation management process all represent control towers. For example, the operating committee is a multi-discipline body that makes balanced decisions necessary for achieving revenue, profit, and return on investment plans. To do so, it pulls levers including capacity, materials, and labor. It leverages information from many sources in order to be properly informed and to make good decisions.

It is important when constructing a control tower to understand the span of control across functions and time. Does the control tower only focus at the operational level and only on the distribution function, or is it more functionally end-to-end and does it also handle tactical and strategic decisions?

From a technology perspective, it is still possible to go to a central control room outfitted with large screens showing the status of all supply chain nodes across the globe, with the ability to zoom in and zoom out as needed. However, in today's distributed and mobile world, such capabilities can be delivered on-the-fly, anywhere, to anyone (with authorized credentials). Thus, control tower technical capabilities are delivered both to central locations and to distributed organizations and individuals. It is also possible to construct a single technical architecture capable of supporting multiple control towers at multiple levels of the enterprise.

Technology Capabilities

This article discusses basic principles behind control towers and is not intended to go into depth on technical (software) capabilities. However, this section summarizes some of the important software capabilities of modern control towers. It is not intended to go into great depth on each capability, nor does it describe how the capabilities work together.

Model of the real world

Making good decisions requires an understanding of the problem. Einstein once said that if he had sixty minutes to save the world, he would spend the first fifty-five minutes defining the problem and then the final five minutes solving the problem. Control towers (indeed, any supply chain decision capability) start with a robust data model that provides a detailed description of what is happening at any point in time. The effectiveness of a control tower is dependent on how well its data model can represent the process to be controlled.

It has become fashionable in recent years to call this a "digital twin." The idea is that the model and its data should be an exact replica of what's happening in the real world at any moment; hence, it's a twin. A common problem with models is what is called "drift." Drift means that the model provides a good representation of reality when it is first used, but the model then "drifts" away from reality over time. Therefore, it is important to have a model that can evolve alongside the real world it represents. This is done by incorporating learning capabilities that provide predictions on what is really happening in the real world and then adjusting the model accordingly.

The technology behind the digital twin is typically an in-memory database (see "In-Memory Computing, SCM Software, and Moonwalking with Einstein"). In-memory databases are used to provide extremely fast response times at scale.

Links to data sources

In order to load the digital twin with data and to keep it updated with what's happening in the real world, it must be linked to data sources. These data include capacities, lead times, orders, deliveries, and shipments. The path between data sources and the digital twin is known as a "data pipeline." In the past the data pipeline was made up of a series of batch data movements enabled through extract, transform, and load (ETL) technologies. Today, the world is increasingly moving towards real-time updates enabled through streaming architectures (see "Data Synchronization in Supply Chains").

On-boarding

The "network control tower" is a type of control tower in which supply chain participants (companies) join a network. The network manages data access and connections. This contrasts with enterprise control towers, which are contained within an enterprise. In the enterprise case, supplier, partner, and customer data flows directly to the enterprise; it is then incorporated into the data model for decision purposes. A key capability of the network control tower is the ability to easily on-board new nodes (companies) into the network.

A plan, or objective function

In order for a control tower to perform its control function, it needs an objective function. Control tower objective functions are entrenched in a plan. This includes detailed plans for sales, distribution, transportation, production, and supply.

Events and codified resolution playbooks

Events represent threats and opportunities. This includes low sales, unplanned orders, cancelled orders, excess inventory, stockouts, low yields, and late shipments. Any model variable that is out of its planned range represents an event that has to be dealt with.

Decision playbooks are codified paths by which to resolve events. While much attention has been paid to AI-supported decision making, there is still a lot of value to be derived from consistent rules-based decision making. In fact, AI support should be layered on a solid foundation of rules-based decision making. The playbook has a simple three level format: deviation->root cause->resolution lever, as shown here (this is further discussed in A Rudder for Course Correction).

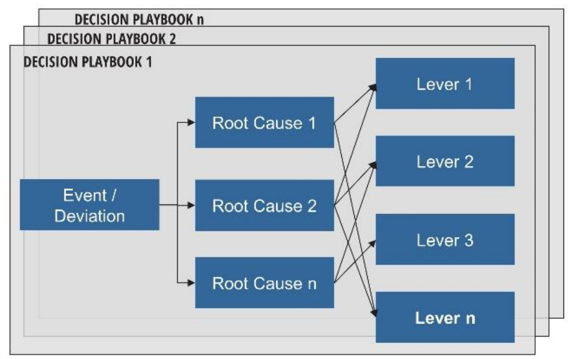

Algorithms and analytics

Control tower and similar software incorporates a variety of algorithms and analytics to support decision making processes. (The terms “algorithm” and “analytic” are often used interchangeably; an analytic is often the result of running an algorithm). Algorithms are basically math that is run against the model in order to make a prediction. For example, the creation of a plan incorporates optimization- and heuristic-based math in order to arrive at a plan that the enterprise can feasibly execute in order to achieve its business goals. In this case, the plan is basically a prediction.

Algorithms are also used on top of decision playbooks. When a path is chosen to resolve an issue, the algorithm is run to determine the overall impact of the chosen path. This is also called a scenario. When operators want to look at many possible resolution paths, many scenarios are run and then compared to determine the best possible resolution path.

Decision making

Control towers are about making decisions – when issues and opportunities arise, analysis is done and decisions are made to resolve the issues and take advantage of the opportunities.

Decisions can be organized into three methods:

- User (manual)
- Algorithm-assisted
- Automatic

Control towers typically employ a combination of these methods. In fact, almost all decisions today have some algorithm assistance and automation associated with them. For example, plan creation has been automated since the dawn of the use of computers for MRP bill-of-materials explosions more than forty years ago. Plan quality become progressively better as more sophisticated algorithms were added.

Likewise, while it is still possible to interactively make manual changes to resolve problems, behind the scenes there are still algorithms being run. For example, a manual change from one source to another is accompanied by algorithms that are immediately run to determine the impact of such a manual move.

Collaboration Framework

Collaboration is essentially joint decision making across multiple parties. In its simplest form, this means that everyone impacted by a decision has visibility of the context of the decision and has input into the decision. Of course, in enterprises, decision processes are typically formalized; thus, collaboration happens in the context of a specific enterprise's decision process.

It's important for a control tower to have robust collaboration capabilities. For example, when an event occurs, it is routed to all those parties to which the event is important, and it is then assigned to someone for resolution. The resolution process may involve considering alternatives, in which the assigned person collaborates back and forth with others on which alternative is best. This process is increasingly automated, with the system automatically generating the scenarios and then communicating the results to various functional experts who then collaborate and consent to the answer.

General purpose enterprise collaboration frameworks have become very popular in the past five years, displacing email as a primary form of communication. Supply chain control towers not only have their own context-specific collaboration frameworks but are also are increasingly integrating with these general enterprise frameworks.

Scenarios

Scenarios are among the most important capabilities of a control tower. The ability to rapidly create, run, and compare scenarios is important to making good decisions in a timely fashion. Scenarios are the means by which the different paths of decision playbooks can be evaluated and chosen. They are also important in evaluating probabilities.

Automation has also now been incorporated into scenarios: when an event occurs, a set of scenarios is automatically created by instantiating the decision playbook for that event. Each scenario represents a playbook path for event resolution.

Learning algorithms
Learning algorithms are now being employed by leading companies to improve model integrity and to eliminate model drift. For example, these companies are employing such algorithms to understand probabilities associated with range-based variables such as transportation lead times and production yields. There is variability in all supply chain processes, so such variables typically operate within a range. Historically, all model variables were treated deterministically, meaning they were single data points; furthermore, they were treated as if they did not change much over time. That is now all changing with the addition of learning algorithms that can constantly run in the background and make change recommendations.

So far, learning algorithms for the decision-making process have been focused on short-term decisions, for example what and how much to order in the next day or two. The challenge with decisions that have longer-term outcomes is defining what goodness looks like in those outcomes, correlating those outcomes to input variables, and gathering data from the future as it happens three months, six months, and sometimes years after the decision. For example, AI's used in college admissions correlate candidate decisions with outcomes five and ten years down the road and beyond. However, in this case, it easy to draw a correlation line between a decision and an outcome, whereas in enterprises, performance in the future may be the result of many factors, including many decisions across many organizations.

Stochastic (probability) analysis
Decisions in supply chains (and generally) are made under circumstances of uncertainty. Demand, in particular, has a high degree of uncertainty. Likewise, many elements of supply have uncertainty associated with them. This uncertainty is mitigated through decision-making agility and intelligent use of buffering. Scenarios are particularly useful in dealing with uncertainty. If scenario capability is robust enough, it can essentially be used to perform probability analysis.

User experience
Control towers must provide a user experience that makes the underlying complex capabilities simple. Since control tower decisions impact multiple organizations, the user experience must also have rich collaboration capabilities across its constituencies.

Final Word

Control towers are an important capability in a world that places high value on precision. Chances are that you already have some level of control tower capability in place, whether within a single function or across multiple functions and multiple decision horizons.

The need for ever higher levels of precision is driven by voracious customer appetite for tailored products with tailored fulfillment. In turn, this is driving tighter tolerances around everything in supply chains: delivery windows, fulfillment times, inventories, warehouse pick times, and production schedules. Control towers are an effective means by which to manage these increasingly tight tolerances. Furthermore, control towers are effective at reducing friction (or mitigating friction) across organizational boundaries, where precision has historically been lost.

While control tower technologies can be complex, the underlying principles behind control towers are simple. These principles are based on closed-loop control thinking, which has been around for decades. Understanding these principles and how to apply them to supply chain problems can reduce the hype associated with control towers and focus more attention on delivering value to customers, employees, and shareholders.

Kelly Thomas is a supply chain management professional and former software executive. He is currently CEO of Worldlocity, where he researches and writes about software, supply chain management, and related topics. He can be reached at *kelly.thomas@worldlocity.com*