The Agile Coaching Company

Agile Strategy Map

A collaborative framework to design, manage and support execution of strategy

Introduction

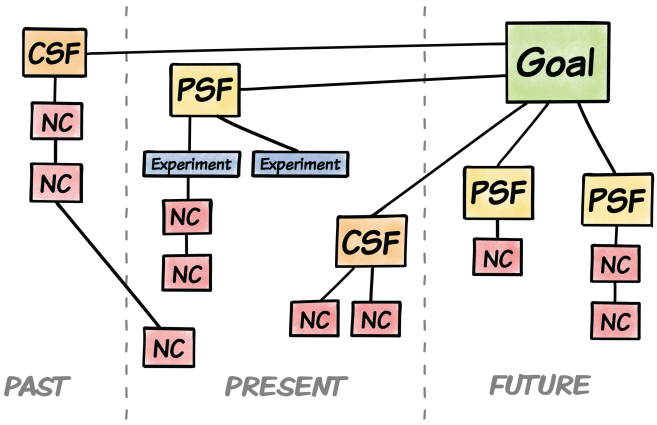

The Agile Strategy Map can be a stand-alone tool for your organization or it can be used in the context of an approach inspired by the principles of ORGANIC agility. In this case, it corresponds to the basic principle of validating changes in small increments. In this document we are going to first present some arguments for the importance of delivering value in small increments to an agile way of working. We will then discuss engaging people in the process of an agile transformation and the risks of pushing change. In the rest of the document we will present instead an alternative way to map and design the desired changes that make the process transparent, incremental, available to everyone, and based on continuous experimentation and adaptation.

Agile teams deliver in small and frequent increments

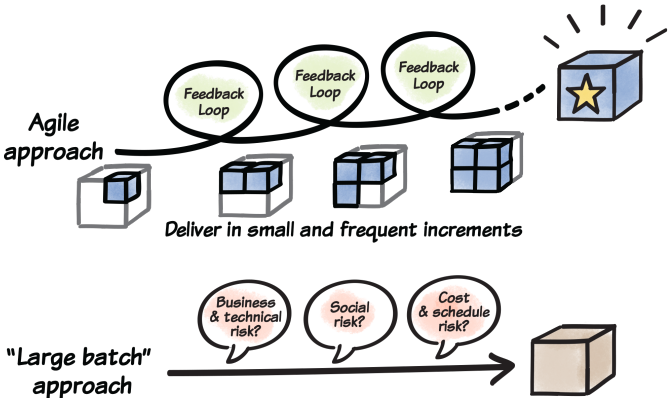

One of the strengths of any team or organization working in an agile way is their capability to deliver value in frequent iterations and in smaller increments. This capability has significant advantages compared to a more common “large batch” approach. First of all, it allows teams to go through the problems they will face from top to bottom and deal with all technical and functional difficulties very early. Moreover, it enables faster feedback loops, which contribute to maintaining focus and directing the team towards what matters to customers and stakeholders.

In a nutshell, this means that working with smaller and frequent releases significantly helps reduce business and technical risk, by delivering what the customers expect and constantly

ensuring that we are using the appropriate technical approach. Moreover, the frequent release of value-focused increments also helps mitigate social risk, by fostering the need for collaboration and trust between team members from the early phases, which avoids stress down the line. Finally, the frequent involvement of customers and stakeholders allows teams to better control both costs and value delivered and enables better expectations management and an agreement on what to invest in next, reducing cost and schedule risk. These benefits apply whether we are developing a new product or introducing an organizational change.

Engaging people with a change initiative

The advantages of an agile approach have often led to Agile being considered as a goal in itself rather than a means of achieving your business goals. Even if this problem is known in theory, in helping many organizations start the journey towards becoming more agile, we have often seen a push towards “making the organization agile” in practice.

When the Agile Manifesto was written, the underlying thought was never about working faster, or being more efficient, or writing more code lines per hour. In fact, the manifesto tells us quite the opposite, by proclaiming the importance of achieving what is called a “sustainable pace”. This is of course relative to the context and the people involved, and it is something to aim for, not a given, or something that can be calculated, exploited, or standardized across an organization. Every team will normally find their own sustainable pace over time, a working speed at which nobody feels stressed, but also not too relaxed. (This speed is, by the way, something you measure, not something you define before you start). The idea of the sustainable pace also has to do with the capability to sustain that pace indefinitely. This last element is quite important, as in most organizations the belief that we can push people to work above their capacity, even if just for a while, seems to be accepted as the norm.

Pushing people in general is not a good choice, in particular when we are dealing with knowledge work. It has been demonstrated1 that when working under stress, our brain capabilities degrade significantly, and a significant reason for people working over capacity is being planned on multiple projects. Under stress, the part of our brain that takes control is the instinctive one, also known as System 1 Thinking2. System 1 is very efficient, but not very

1 There is a lot research pointing to the fact that multitasking in particular, but stress in general, affect people’s IQ negatively, in the moment as well as in the long run, by hindering their capability to be creative and focus. This article provides a good overview and points to various other researchers: https://appliedpsychologydegree.usc.edu/blog/to-multitask-or-not-to-multitask/

2 Thinking, Fast and Slow is a best-selling book published in 2011 by Daniel Kahneman, Nobel Prize in Economics laureate. In 2012 it won the National Academies Communication Award for creative works that help the public understanding of behavioral science, engineering and medicine. More can be found here: https://en.wikipedia.org/wiki/Thinking,_Fast_and_Slow

flexible, and normally resorts to known patterns and mechanical reactions to known “threats”. When we need to be creative, we need the slower part of our brain to work for us, which is part of System 2. This is the part that is in charge of our analytical thinking, of our dreaming and creative ideas, the one that allows us to invent new things, instead of efficiently processing known ones.

If we understand the deep connection between the constraints and structures that we build into an environment, and the behaviors these structures produce - in terms of human reactions to stressors - then we will want to make sure that changes happen without too much pressure. Even though we all know the benefits of working without pressure, or after a good night’s sleep, as soon as we are again in an environment where we are pushed to do things fast, we inevitably fall back into compliance mode and retreat into behaviors and habits that are typical of our traditional way of working.

Let’s assume for a moment that your organization is willing to go through a change initiative, and the benefits are perceived as so valuable that nobody doubts that it will be worth pursuing. Even in this hypothetical and very unlikely thought experiment, how much could we resist pushing people through it, instead of allowing people to internalize the change at their own speed and co-evolve with the system? Since we know the answer, we feel that allowing people to move towards accepting the change, as a first step, would be a waste of time.

This is the moment when, instead of being supportive and respectful of everyone’s need to understand and adapt, we inadvertently begin generating what we call “resistance to change”. Continuing with our thought experiment, what if we were able to share a common and measurable goal that everyone would understand and would be willing to pursue over time? What if, instead of telling people what to do, we would allow them to try safe-to-fail experiments geared toward achieving that goal? Based on the results of these experiments, we would encourage people to share their approaches and replicate these conditions across the whole organization. As crazy as it sounds, you will be surprised about how much faster and more sustainably these changes will grow within your organization.

Back to reality, we all know that finding ourselves in a situation in which everyone within an organization, no matter the size, agrees to change or to a common approach to change is very rare. This next section is about identifying ways to create the sense of urgency that will allow us to initiate change. Instead of making it explicit, or a mandate, we will explore ways in which it can emerge by analyzing common needs or dissatisfactions. Any change requires a motive, and this motive needs to be internalized by as many people as possible in order to make the change not only accepted, but also thoroughly lived. In this way we can also avoid the risk of local optimization, which can be a result of naive efforts to increase local efficiency, rather than

focusing on client value delivery by taking the client or market perspective. We want to find ways in which we can engage step by step the whole organization, and align it behind a common agreed direction.

Managing uncertainty through experimentation

The idea of experimenting is quite interesting and engaging for most of us coaches, but it is more often than not a tough sell for leadership teams within organizations. It is really hard to convince hard driven managers, with targets and deadlines, to experiment on something without knowing what the outcomes will be. It ultimately boils down to one of the hardest things to accept: *uncertainty*. When talking to top managers and C-Level people instead, the level of acceptance is much higher due to their natural predisposition to keep multiple options open. Paradoxically, it is easier to sell top managers a portfolio of safe-to-fail experiments with different probabilities of success and a clear strategy to manage the risk, than to discuss starting or not a single experiment.

A not so new approach to risk

Nowadays it is very hard to base our risk management strategy on robustness, in this case predicting and preventing failure. Risk used to be visualized as a bell curve, which makes its management complicated and slow, by focusing on what lies inside the bell curve alone. Rare, one-off events (Black Swans, in Nassim Nicholas Taleb's terminology) are impossible to predict that way, and are therefore especially destructive — they have a greater impact. If we rely on what we think we "know", what will potentially hit us the most is what we actually don't know, or what we think we know just "ain't" so. By accepting the possibility of failure, an organization can instead orient its strategy towards early detection, fast recovery and fast exploitation. Statistical techniques are still valuable for ***probable*** events, and simulations and scenario planning allow us to gain some clarity in the realm of **possible** events, but when it comes to the enormous number of **plausible** events, we need to use abductive logic to draw connections between multiple events and seeking coherent (but not necessarily true at this point) explanations.

In order to broaden the spectrum of options available and to be able to discover what might help and what not, we need to approach risk management as a whole in a more experimental way. Focusing on probability alone will not work; we will have to learn what is possible and what is plausible. We can formulate hypotheses based on intuition, but those hypotheses will not be used to make decisions, but rather to start multiple parallel experiments to validate those hypotheses. This approach requires us to be open to observing and listening to everything that happens, especially if it is unexpected.

The comprehensive framework of the Agile Strategy Map has proven to serve as a good catalyst for change and coordinating multiple, parallel experiments.

The Agile Strategy Map™ and its background

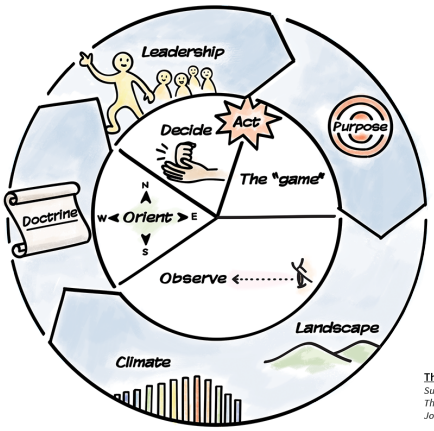

The Strategy Cycle
Sun Tzu's five factors
The two types of why
John Boyd's OODA loop

Now we have learned something about incremental and iterative change, we have learned something about the importance of initiating change from the customer's perspective, and about engaging people within the organization from the early stages, so that the need for change will emerge and will start to be "pulled" by volunteers who feel close to the improvements and feel they can contribute. We have also described the possibility of evolving incrementally, and using an experimental approach to validate hypotheses before getting started. These options are all built into the Agile Strategy Map. This framework has been developed through our experiences with clients and the help of many coaches who contributed over time to refine and improve its usability.

The Agile Strategy Map's origins are rooted in Eliyahu Goldratt's "Strategy and Tactics Tree"[3], a thinking process codified in his Theory of Constraints. This provides a model for aligning and synchronizing continuous improvement. The Agile Strategy Map tool evolved into a framework that can be used in multiple circumstances: it helps with maintaining coherence towards a common goal, aligns everyone on the current state of affairs, and allows us to straightforwardly track dependencies. It also merges strategic priorities with tactical and operational needs, allowing for a more focused approach.

Additionally, the work of Peter Senge, presented in his book *The Fifth Discipline, the Art and Practice of a Learning Organization*, has been a significant influence. Validating change in small increments is essentially about building a culture and discipline of learning, rather than simply defining a plan that we presume will result in achieving our Goal. To quote Senge, *"In the long run, the only sustainable source of competitive advantage is your organization's ability to learn faster than its competition" ... "If there is one single thing a learning organization does well, it is helping people embrace change"*. Furthermore, Senge's "Wheel of Team Learning" provides a simple way to consider the process of collective learning and reflects the set-up and running of the Agile Strategy Map: we identify shared needs and goals (Shared Meaning, in Senge's terms),

[3] If you are interested in knowing more about Eliyahu Goldratt's "Strategy and Tactics Tree" you can have a look here: https://www.toc-goldratt.com/tocweekly/2011/06/gst_a_step-by-step_guide_for_change/

co-create the Agile Strategy Map (Joint Planning), agree on validating through collaborative actions (Coordinated Action) and make it transparent for all to see (Public Reflection).

Wardley speaks of real maps not only being **visual** and **context-specific**, but also showing **positioning** in relation to an **anchor** and **movement**. Many things that we use and call maps in the workplace are not maps at all, lacking at least one of those elements. If they don't include an anchor (that would be the North on a geographical map) to support clear direction and positioning to show us where we are in relation to other elements on the map, then how can we use the map to orientate ourselves? Wardley has combined the thinking of OODA loops (the decision-making cycle of observe, orient, decide, and act) from military strategist John Boyd and "The Art of War" from Chinese general Sun Tzu (VI-V century b.C.) to create a basic cycle for thinking about strategy.

Using this cycle we begin to see that the strategy changes based on the needs and the maturity of the market and where we want to go next. This has of course a significant impact on the way organizations are structured and operate.

To sum up, based on Wardley's insights, good strategy tools are:

- **Visual** → easy and intuitive
- **Context specific** → relevant to the context (different parts of the business might have a different strategy or different products and it is important to know what is universally applicable and what is not)
- **Positional** → displaying connections and current state to help navigate the map
- **Anchor** → acts as a reference for direction
- **Movement** → suggested changes and where are we going / where have we been

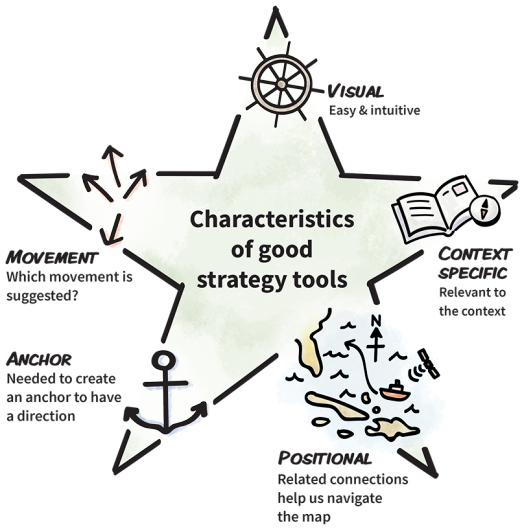

Elements of an Agile Strategy Map™

The Goal

The Agile Strategy Map is a real visual map that includes the elements of anchor, position, and movement. In principle it is a way to visualize a goal, as well as the success factors and dependencies that are relevant to moving in the right direction. The **anchor** of an Agile Strategy Map is a Goal, which can be expressed as a specific target, measurable and timed,

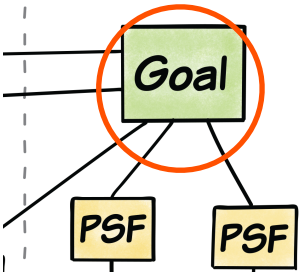

or simply as a direction towards which to move. It represents the business goal and creates focus for the strategy, orienting all other elements. Since the Map is also **context-specific**, the Goal will need to fit the context of where the organization is and where it intends to go.

It is also very important for the Goal to be outcome-based, not output-based, which means that it must be connected to creating value to users, customers, and/or employees. Stakeholder value is a product of the fact that customers and users are satisfied and employees are engaged.

If the Goal is a specific target, it is possible to use different techniques to formulate it, such as the SMART checklist[4]. An example of such a Goal can be: *"Strengthening our position in mobile services by increasing the number of annual mobile service's customers by 20% over the next 12 months"*. Alternatively, we have in the past successfully used the "Remember the future"[5] technique, which is based on numerous studies in cognitive psychology investigating how we think about the future: *"Imagine that you fall asleep now and wake up in 12-18 months. What key changes do you see around that make you happy?"* These kind of questions generate more richly detailed and sensible goals, because it is easier to understand and describe a future event in the past tense than a possible future event, even if neither has occurred. By thinking of a future event as one that has already occurred, we also pave the way for imagining possible factors that will enable or accelerate generating the event.

In the case of the Goal as a direction, according to complexity thinking, it can be expressed and measured in terms of Vector Tracking (as direction and speed of change). The target for the organization can then be the direction and speed of change. For instance: *"We want to increase our customer satisfaction 20% faster than it is currently growing, so we will be outpace out competition and increase our market share significantly"*.

Even if we have defined the Goal, we want a way of reminding ourselves that we should challenge what we are trying to achieve as often as possible, because reality and conditions around us change very quickly. The volatility we are dealing with nowadays is such that it is very risky to base medium- and long-term plans on current situational analysis without planning for continuous adaptations. The Goal itself might be discovered over time, or at least refined, if not reshaped, by integrating new insights. The Agile Strategy Map process is supportive of continuous adaptations and injections of new insights, allowing us to maintain coherence and transparency at the same time.

---

[4] You can read more about all different dialects and variations of the SMART check-list here: https://en.wikipedia.org/wiki/SMART_criteria

[5] This is one of the most common Innovation Games in the agile community. You can find more information on how to use it here: https://www.innovationgames.com/remember-the-future/

Exploring the existing Landscape

Virtually every organization has some kind of strategy, or at least has a plan to get to some definition of success. We recognize that most organizations have achieved some level of success in the past, whether intentionally or just out of luck. Even when these strategies are very static and expose an underlying linear and mechanical thinking, it is important to show appreciation for what the organization has achieved, and identify what has helped the company be successful in the past.

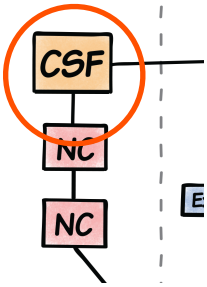

Confirmed Success Factors (CSFs) are an expression of the successful factors that have led company to its current state and will provide a picture of the past Landscape and Patterns (to refer back to the concept of the map). These - in line with the ORGANIC metaphor - could be considered like an organism's acquired capabilities, which became part of its DNA as a result of an evolutionary process. The CSFs might be in the form of processes, rules, policies, constraints, approaches, in short everything that is established as a way of working within the organization, as well as established value propositions to existing customers. All these things, learned over time and validated, are assets to the organization, and are probably responsible for a significant part of the overall revenue. Given the defined Goal, we may be able to identify a subset of CSFs that will be enablers for achieving the goal. We want to be clear about focusing on those that we believe to be relevant to the Goal and its specific context. This might seem like a hard decision, but if you want to achieve success you need to focus on what is most important to you and relevant to your business.

The term "Confirmed Success Factor" emphasizes that we have attained some knowledge and that this lesson has been retained and consolidated into an organizational asset. A CSF is, then, always in the *Obvious* or *Complicated* domain according to the Cynefin framework.

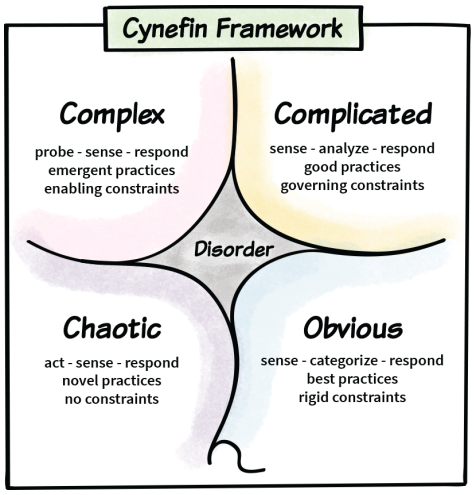

Cynefin provides a way to understand the context we are in and the most appropriate patterns and constraints to be used in it. It assists the decision-making process by identifying different approaches to situational analysis and the decision itself, which depend on the **domain** we are in at any moment in time. Cynefin defines five different domains, which are divided in three categories: *Ordered Domains*, *Unordered Domains*, and *Disorder*. By definition, the latter is the domain in which we are, when we aren't able to determine which domain we are actually in. The *Ordered Domains* are the domains of causality, where the

connection between cause and effect is clear. There are two domains in this category:

- the *Obvious* domain is characterized by known qualities and is governed by rigid constraints (or policies). It is the domain in which Best Practices can be defined. According to the definition of this domain, the situational analysis is self-evident, while the decision-making process is supported by rules, checklists, and defined processes. To ensure that these are followed, regular reviews are scheduled.
- the *Complicated* domain consists of known unknowns. There are still Governing Constraints and Good Practices, which can be documented and are usable by experts to increase chances of success. Once they are documented, we need to link them to daily work and processes, and maybe, in time, turn them into Best Practices that can be applied by anyone. Till then only experts can use the Good Practices effectively. In this domain the situational analysis needs to be done by experts, who will identify all possible options or alternatives, and will present these for a decision.

A Confirmed Success Factor may be expressed in the following form:

WE LEARNED THAT <SOMETHING> **CAN BE LEVERAGED TO ACHIEVE** <GOAL OR OBJECTIVE> **AND WE CAN MEASURE IT WITH** <LIST OF METRICS>

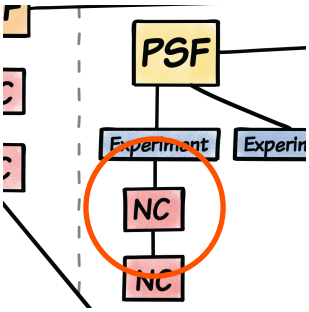

As mentioned previously, it represents an acquired capability for the organization that can act as an enabler towards achieving the goal. The fact that the CSF is achieved and known doesn't mean that we won't have to do anything about it. On the contrary, a CSF is like a lever that we can use to enable our organization to achieve success and needs to be oiled and maintained, or it will decay and loose relevance. To maintain and continuously evolve a CSF we require at least one **Necessary Condition**. This can act as an anticipatory trigger, reacting to or prompting specific events/needs, for example periodically reviewing a policy to check how it's performing against some Key Performance Indicators (KPIs). We can create triggers in the form of Necessary Conditions, connected to the KPIs or to a specific moment in time. The dimension of time is also integrated into the Strategy Map, so if the NCs are connected to a date (likely at some point in the future), they should be placed in the Future column of the strategy map, while still being linked to the relevant CSF. If we are unable to define what is necessary for maintaining the Confirmed Success Factors, this may be a sign that they have not actually been confirmed/validated or perhaps that they are no longer relevant.

Define hypotheses to test explicitly

Going back to the Cynefin framework, and looking deeper into the *Unordered* Domains, or the domains in which it is impossible to determine causality without uncertainty, we can recognize two different domains: *Complex* and *Chaotic*. In the *Complex Domain,* because we do not yet know what we don't know, the path towards validating a Goal is never straightforward. Most of the time we have to understand and analyze our hypotheses and challenge our assumptions in order to figure out our next move. For this reason, the next step in the creation of an Agile Strategy Map is the definition of those hypotheses that might help us move towards the Goal. These hypotheses need to be made explicit, so that through transparency dependencies can be made visible. The primary purpose of declaring explicitly what could be helpful towards achieving the goal, is to identify changes or adaptations that can be used to our advantage. Ideally we would want to have many alternative hypotheses available, and we shouldn't discard them right away. At this level a good set of 10 to 14 different hypotheses, would provide enough options to explore and avoid focusing only on the obvious ones. Hypotheses can be naive, or even completely stretched: as long as they are plausible and coherent, they are good. These hypotheses are captured using **Potential Success Factors (PSFs)**. The name is a reminder that they are still to be validated.

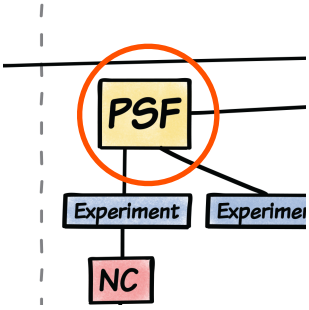

A Potential Success Factor is expressed in the following form:

BY <DOING SOMETHING> **WE EXPECT** <SOMETHING TO HAPPEN> **THAT SHOULD SUPPORT ACHIEVING** <GOAL OR OBJECTIVE>

Given the example of Goal: *"Increase the number of annual mobile service's customers by 20%",* an example of a PSF can be: *"By creating new free services, we expect to attract more people to our platform, that should support achieving an increase in the number of paying customers by 20%"*. PSFs are designed to be validated or invalidated through rapid experimentation. After they are validated, they will provide more insight into our strategy and increase or decrease the level of confidence in moving forward in one direction or another. If we feel confident about a PSF then it will eventually be converted into a Confirmed Success Factor (CSF). Once we have defined the PSFs, we visualize them underneath the goal to make them transparent and take full advantage of the visual capabilities of the tool. Since it is important to base decision-making on context, we have to make explicit which kind of hypothesis is described in each PSF: The Potential Success Factors either represent known unknowns (which then means we are in the *Complicated* domain), or unknown unknowns (in which case we are in a *Complex* domain).

Decide what to focus on

As we said at the beginning, if we want to be able to focus on small validated changes, we must decide which PSFs we want to work on first. Contrary to a more traditional way, we don't want to actually prioritize the PSFs but rather make them smaller so that we validate their impact on the goal faster and more effectively. Every Success Factor (PSF or CSF) should have a ***Champion***, who will work to build a cohort that can collaborate and focus on moving the PSF forward, and who will remain the Champion if the PSF becomes a CSF. The cohort is what we call an ***Improvement Squad***, as its objective is to improve the organization and the work of everyone involved, not to mention the results, by exploiting new capabilities or leveraging existing ones.

Identify Necessary Conditions to validate the hypothesis

We need to find ways to validate our hypotheses as fast as possible, empirically, and without relying on assumptions that ultimately increase risk. This can be achieved by designing small, safe-to-fail experiments. Before getting there though, we need to identify what it is necessary in order to be able to define such experiments. What do we need to have in place or deal with in order to be able to validate the hypothesis? These may be things we need to change or implement, or they may be constraints that we must address in some way. These "necessities" will also be captured using Necessary Conditions (NCs) which should also highlight (in the Experiment Canvas capability of the Agile Strategy Map) what could go wrong if they aren't fulfilled. This helps prioritizing and identifying dependencies. Once all the NCs have been fulfilled, we should be able to define one or more experiment(s).

A Necessary Condition may be expressed in the following form:

WE NEED TO <....> OTHERWISE <...>.

Given the example of the PSF above, an example of NC can be: *"We need to create at least one additional free service in order to measure increased subscriptions, otherwise we won't be able to understand the impact", or "We need to measure existing conversion rates, otherwise we won't be able to set an appropriate target and measure the increased conversion because of free services".*

In short, the Necessary Conditions will bring the strategy to a tactical level and allow operational work to start. They help in either validating a PSF, in planning the roll-out of a newly identified CSF, or in structuring the management of an existing CSF.

Relationships between Necessary Conditions and PSFs/CSFs/Experiment Canvas give different meaning to a NC depending on where it is visualized on the Strategy Map. Here is a summary table defining the meaning of each specific relationship.

Strategy Map element	Position on the Map	Meaning of NC
Possible Success Factor (PSF)	Present/Validation	What do we need in order to create an experiment to validate this PSF?
Experiment Canvas (EXP)	Present/Validation	What do we need in order to be able to start this experiment?
Confirmed Success Factor (CSF)	Present/Validation	What do we need in order to make this an asset for the organization? Which training, changes, automation, policies…
	Past/Confirmed	What do we need in order to monitor, measure, and maintain this CSF? Do we need to create any anticipatory triggers?

Tips

Tip 1 If we see the pattern of a NC being applicable across multiple PSFs, we can have a reasonably high level of confidence that we will get a high return on effort if we manage to effectively address that condition. In this case it doesn't matter in which order we tackle the Success Factors, because the first one will need to address the NC, and therefore the second one will need less work to move forward.

Tip 2 What level of granularity makes sense in the Agile Strategy Map? This is a very common question and there isn't a clear answer. However, there are some considerations that can help answer the question on a contextual basis:

- If the level of granularity of Success Factors is too low, we might end up with uncertainty regarding whether we are dealing with a CSF or a PSF, because it might be both. To make it clearer, here is an example Success factor from a Car manufacturer: *"Cars with low emissions are easier to sell". "We learned from experience that this is true", "We have data backing this fact", "We know how to reduce emissions today"*. The above three statements are verified. *"We need to keep researching how to make cars more efficient in the future".* This is not verified and requires experimentation. We don't know yet the how but we know that this might be important for continuing to sell. So how would we represent this on the Agile Strategy Map?

Option A: This could be a Confirmed Success Factor (known knowns) with a Necessary condition that states the need to review the technology every three months, which in turn could start a research initiative on the portfolio.

Option B: Break the CSF down into multiple and more specific CSFs, which highlight what is currently known, and a PSF for the hypothesis. Option A is preferable because this stays at a level where it avoids getting bogged down in technical details. It would also allow the same Champion and Improvement Squad to maintain control of what is known, but also keep growing the acquired knowledge so that it never becomes obsolete and irrelevant.

Create safe-to-fail experiments

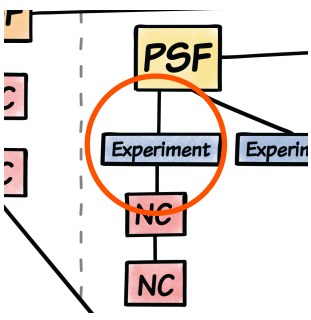

As soon as you have identified which are the Potential Success Factors you want to focus on, you pull those from the Future/Potential position to the Present/On going position of the Agile Strategy Map and start creating experiments to validate those hypotheses. To get quick feedback and make decisions, the recommended duration of the experiments is 4 to 12 weeks. If we go back to Cynefin and complexity thinking, we can see that experiments in the complicated domain are meant to evaluate possible

options, while experiments in the complex domain are meant to let new options emerge. Therefore, for the complicated domain, we run one experiment and validate it. When dealing with situations in the complex domain, we suggest running multiple parallel experiments, as the context in which the experiments are executed might change quite rapidly. By having multiple parallel experiments, we will be able to recognize recurring pattern(s) across those experiments, identify the possible catalysts that sustain those patterns, and finally validate that what we have identified are actual catalysts by testing those on all the experiments in parallel. This type of approach isn't possible when running a single experiment. The quality of the situational analysis will also be greatly amplified by having multiple different datasets. The recurring patterns might lead to options (the identified catalysts) for which we want to define additional experiments, now in the complicated domain, in order to evaluate the most appropriate one(s). We use an Experiment Canvas, integrated in the Agile Strategy Map framework, to help articulate what are the things we need to know and measure when running an experiment.

Decide which experiments to start

With small validated changes in mind, we probably won't start running all experiments at once, but decide which are the most appropriate to run first, pull them into the central column of the Agile Strategy Map, which is the present/validation column, and place the others in the future/ potential column. The faster you can validate or invalidate your hypotheses, the earlier you will have an understanding of how to develop your strategy. In this phase more Necessary Conditions might emerge as preconditions for experimentation, depending on the level of complexity you are dealing with. As soon as you have fulfilled all the NCs required for the current experiments, there should be no further delays. Every experiment should have predefined success and failure conditions, as well as amplifying and dampening actions[6]. Identifying those before you start helps you make validated choices without unnecessary interference and track progress over time[7]. It will also help to empower the team which will run the experiment, by providing clear boundaries and suggestions on what to do when specific success or failure conditions are met.

Reporting on experiments' results

When running an experiment we don't want to wait 8 - 12 weeks to assess and communicate if it was a success or a failure. You can instead visualize the real-time status for each experiment: Each time the experiment team decides that one of the success conditions in the experiment is fulfilled, they are empowered to tick it off. The same applies to failure conditions. This is mapped

[6] to learn more see http://cognitive-edge.com/methods/safe-to-fail-probes/

[7] agile42 in collaboration with Cognitive Edge has developed a tool called Organizational Scan, which allows the user to relate decision-making speed, emotional responses and other indicators to experiments connected to the Agile Strategy Map. The tool permits real-time monitoring of each experiment's progress, and trigger-based intervention, to enable the next generation of Strategic Decision Making platforms.

by a line moving upwards or downwards to indicate the level of success/failure over time, which can increase confidence in making decisions on how to move forward.

Roles

The following roles connected to the Agile Strategy Map are not mandatory, but we have found them critical for the success of any strategic change:

- **Sponsor**: The Sponsor is usually at the C-level or an executive leader. The Agile Strategy Map Sponsor is the person responsible for the budget the organization needs to invest in change. She advocates for the overall strategy and supports the Champions and the Improvement Squads in their initiatives. In some cases, depending on the size and complexity of the organization, we may form a Guiding Coalition of leaders who are responsible for dealing with large organizational constraints and insuring overall strategic coherence.
- **Champions:** Each PSF or CSF has a Champion who advocates for it and ensures that it is getting the proper attention. They are usually a senior leader or an opinion maker, who can influence others and help create the proper environments for feedback and learning. Champions will recruit an Improvement Squad as they feel appropriate. All Champions together will form a team to collaborate on the overall strategy map, focusing on individual Success Factors as well as overall map design and movement.
- **Improvement squad:** In most cases, the Success Factor Champion will need help to define and accomplish the NCs, design any experiments needed, or deal with constraints or prerequisites for an experiment to start. Each Champion should recruit an Improvement Squad, a cohort of people who can contribute to moving forward with the necessary actions. Notice that it is a Squad and not a Team because it will change composition over time, depending on what is necessary for supporting the change.
- **Experiment team:** Once an Experiment Canvas is designed, a volunteer-based experiment team is formed to carry out the experiment. This team remains stable for the whole duration of the experiment. While they run the experiment, the Improvement Squad supports the process and monitors the outcome. The same experiment team may be able to run more than one experiment as long as they aren’t dependent on each other.

Evaluate and Validate

Collect data at regular intervals

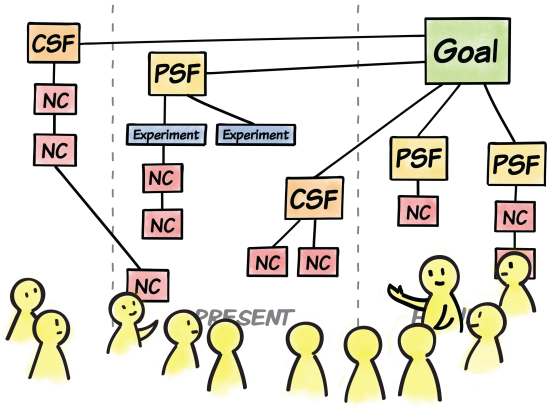

Make sure the Strategy Map is visible to the whole organization and set up a system so that everyone can contribute. There are multiple ways to leverage the collective intelligence and cognitive diversity in your organization. For instance, create a straightforward way for anyone to give feedback on the strategy in terms of Goals, PSFs/CSFs, and NCs. The *Improvement Squad* discussed in "Roles" is an additional way to involve more people. They can visualize the activities related to the different NCs on a Tactical Board, which is both a way to move from strategy to operations and a very powerful information radiator.

Once the experiments have started, you should be able to collect up-to-date metrics regularly. This can happen at very fast intervals, or even in cycles of 1 to 2 weeks. The data should help us understand in which direction and at which speed the experiments are moving (Vector Tracking, as described above), which should allow us to make decisions faster.

In complex environments we have multiple safe-to-fail experiments/options for each success factor. Here, we are trying to understand what patterns emerge, so that we can start amplifying the good (those that give us the results we are looking for) and dampening the bad. Occasionally we discover unintended consequences or hidden patterns that impact parts of the organization or factors that we did not consider. We could end up solving additional problems in this way.

In complicated environments we gather data and evaluate the options. We can then decide if the Potential Success Factor can become a Confirmed Success Factor and how to close the feedback loop to check on the necessary conditions.

Observe the projects interfering as little as possible

We define amplifying actions and dampening actions before the experiment starts. Note that some experiments might be designed to fail, so in that case the "success conditions" will be about failing. The creation of these conditions and actions provides a set of enabling constraints with triggers to action, which helps create a safe-to-fail environment for the experiment team.

Validate the results and learnings

While experiments are running - particularly in the *Complex* Domain of Cynefin - we have to constantly monitor the emerging patterns. To be sure they are actual patterns, we need to evaluate their stability and validate their repeatability by identifying which enabling constraints can reproduce them. These constraints can take the form of *catalysts,* which can both amplify the effects of positive patterns, as well as dampen the effects of negative ones. Will these *catalysts* help to reproduce the positive effects we have observed during the experiment? How could we transfer those learnings and benefit to the organization as a whole? The answer to these questions will help us make decisions about whether to roll out the learnings or not. Remember that we are talking about a *Success Factor*, which should be leveraged to achieve our goal, so if we are unsure about it, then there is no benefit to rolling it out.

Engage with all relevant stakeholders

Engage with all relevant stakeholders and parties in the organization to initially set up the Agile Strategy Map and to understand the implications of a roll-out. Make sure all necessary preparation is complete before roll-out, so that the transition to the new system is as quick as possible. Use the stakeholders to support the transition and engage with all involved to increase acceptance and reduce resistance.

Roll out the change

By supporting everyone involved, finding out fast what works and what doesn't, and providing support where problems arise, you will make your roll-out smoother and more effective. In this phase it is very important to handle all impediments promptly by ensuring through frequent meetings that they are removed as fast as possible to maintain momentum.

After adopting the Agile Strategy Map with dozens of clients, we came to appreciate it also as a powerful Enterprise and Leadership Coaching Tool: the outcome is important, but the conversation is even more important. The impact in terms of sense of ownership and momentum determined by leaders co-creating and collaborating around a common goal greatly increases focus on the business goals, and offers unique opportunities to coach the leadership team towards becoming a more resilient organization.