Crossing the chasm:

From campaigns to always-on marketing

A PEGA WHITEPAPER

By Rob Walker and Matt Nolan

PEGA

Build for Change®

The dire state of marketing

With the emergence of artificial intelligence (AI), traditional marketing is wading into troubled waters – and practitioners don't need a data scientist to feel a major shift coming. We continue to layer consumers with messaging, as we've done for years – strategic campaigns and segments, retargeting, real-time media buys, mobile push, personalized email – but nothing moves the needle like it did a few years ago.

We're in a tough spot: The impact of our initiatives is degrading, and at the same time we're feeling tremendous pressure from new competition. Adjacent companies are moving sideways and vying for our space in the market, and it's adapt or perish, so we constantly look to do more, talk more, and engage in more places. But, somehow, we end up with lower response in return.

In a sector with more than 5000 unique marketing solutions[1], the average campaign response rate is less than 1 percent[2]. That means on average, 99 percent of the messages we send are largely irrelevant.

With response this bad, we think the only way to recover is to increase the customer assault even more: bigger audiences, more campaigns, more pushing. But this has the opposite effect; we end up teaching our prospects to ignore us – that our messages are irrelevant, and aren't worth their time, energy, or focus.

It's a race to the bottom...with everybody losing:

- The customer becomes wary of all marketing, and is deprived of relevant offers.
- The company sees decreasing revenues and fewer promoters.
- The market is threatened with incremental regulation – designed to protect consumers from spam – making it even harder to engage responders.

There's one question every company needs to ask itself – what are we actually trying to accomplish with our marketing? Because the answer isn't to run campaigns. Campaigns aren't the end-goal, just a means to an end. And despite how marketers are driven to behave, the goal isn't simply to sell products, either – that's short-sighted. Instead, the goal must be to increase revenue and profit for the company, as a whole, and, ideally, to increase customer satisfaction at the same time, increasing future revenues.

So... how is it possible that the means got so thoroughly confused with the end?

1 Chief Marketing Technologist Blog, Scott Brinker (2017), http://chiefmartec.com/2017/05/marketing-techniology-landscape-supergraphic-2017/

2 Direct Marketing Association, 2017 Response Rate Report (2017) https://thedma.org/marketing-insights/response-rate-report/

Campaigns were a mechanism popularized in the 80s and 90s, designed to distribute an overall company revenue goal across business lines and products. The idea was that we'd campaign for those products to meet a revenue objective, and call it a success when we sold a certain amount. But at that point, activity masquerades as progress, and we often forget to ask – have we actually improved company-level revenue? Did our product-push come at the expense of a greater opportunity? Did our attempts to sell actually compromise future revenues, customer satisfaction, retention, and so on, or did they create higher risk exposure?

In many cases we're doing more harm than good. With a poor understanding of the end-goal, each business unit squabbling over its share of the pie, and a fragmented view of the customer, we consistently compromise our own successes. Our messages interrupt and distract, rather than add value. By refusing to look at the big picture for our organization and its sustainability, we end up accelerating the downward spiral.

While this "race to the bottom" is easy to see once called out, it's more difficult to actually change – factors like corporate culture, process, technology alignment, and a segment or campaign-focused mindset are all barriers. However, recent order-of-magnitude improvements in real-time interaction and AI capabilities have eliminated technology roadblocks, and the benefits of "doing it right" are absolutely massive – increased engagement, response, customer satisfaction, retention, sales conversion, revenue, and profit margin, to name just a few. The concept of one-to-one marketing was popularized more than 20 years ago,[3] but the technology was never available to actualize the vision. Now, with the advent of real-time AI, and the undisputable benefits to both business outcomes and customer experiences, the only remaining challenges are around culture and inertia.

[3] The One to One Future, Peppers & Rogers (1996)

The three things we absolutely must change

To make the benefits a reality, the following innovations need to happen within the organization - both technically, and culturally:

1. The move from one-to-many segmentation, to a one-to-one customer approach.
2. The move from channel and product-driven targeting, to a central decision authority.
3. The move from scheduled push campaigns, to an always-on model that continuously engages customers during their "moments of need."

These three innovations are not only possible, but early-adopters have already implemented them, and seen massive gains – demonstrating KPI lift across the board, at the same time their status-quo peers were losing ground in the same markets. Given the state of marketing in general, it's clear that we have to change how we think about engaging customers, and be willing to adapt – to invest in an approach that will develop relationships, instead of eroding them.

Next, we'll walk through each of the three requirements, one-by-one, explaining:

1. Why they're critical.
2. How you'll need to change.

Given the state of marketing in general, it's clear that we have to change how we think about engaging customers

Moving from segments to a one-to-one customer approach

The challenges of segment-based targeting

Imagine that a bank wants to sell their palladium credit card – they'll use a campaign to make that happen. With a goal of selling X additional cards, as marketers we'll identify segments to target with the palladium campaign – with the first level of segmentation looking something like this:

- **Age>45 Years; Location=East Coast Only; Income>$150K**

To try and get more personalized, we'll add more criteria, a few levels deeper – looking in the above group for everyone who visited the website, and viewed pages in the Cards section.

- **WebVisit=True; CardsPageVisit=True**

And in many cases, we'd also include attributes about responses to previous offers, eligibility and suitability criteria for the card (e.g. a credit score, disposable income), and their propensity to respond to a similar offer (derived from a previous campaign). If we're marketing in digital paid channels, we'd also include large amounts of (questionable quality) third party data attributes. Altogether, this would read something like:

- **Age>45 Years; Location=East Coast Only; Income>$150K**
- **WebVisit=True; CardsPageVisit=True**
- **CreditScore>500; FrequentTravel=True; PalladiumPropensity>50%**

This is slightly oversimplified, but these are the basics of building an audience with segmentation – a process which introduces severe challenges, such as:

- **Compromised relevance –** Segmentation is a filtering approach, a method that routinely mixes eligibility, suitability, urgency, preference, and relevance criteria – at the cost of precision and transparency. Even when advanced predictive analytics are used as criteria to improve precision, the segmentation scheme is only trying to find an audience for a specific offer. This alone hugely compromises relevance. Relevance is easily the most critical factor in marketing, and it's too complex for humans to solve without using hardcore math, but that's exactly what segmentation encourages us to do.
- **Unscaled labor –** Micro-segmentation limits re-use, transparency, and scale. We know broad-spectrum segmentation isn't that effective, so we keep adding more and more criteria to generate "micro-segments," which supposedly improve our targeting (see above for why that doesn't work). But the more specific a segment gets, the less transparent it is, and the harder it is to actually re-use. Regardless of how they perform, many segments are mostly one-and-done, only used once, then sit on the shelf. Without disciplined documentation and footnoting, after a few weeks, nobody remembers what the segmentation model was even designed to do, much less why those criteria were used. The organization ends up allocating a tremendous amount of resources to build new segmentation models for each initiative, even though they're not very effective.
- **Collisions and conflicts –** Segments often collide and conflict, at the expense of the customer. There's never any guarantee that customers in a campaign segment aren't part of another segment, or concurrent campaign – this happens every day. We usually solve this by referencing one segmentation scheme within another (exclusions), or using 'campaign optimizers' to balance across campaigns. Whatever method we use, things get complicated and murky, very fast. We end up doing whatever will maximize short-term revenue – at the expense of a healthier, more sustainable customer relationships. This kind of problem that needs to be tackled up front, with a strategic approach – not solved tactically, after the majority of the work has already been done.

The traditional campaign model

Regardless of the issues with segmentation in general, we'd push forward – and this is how a traditional marketing campaign works, essentially:

1. **Set the objective –** First we define a campaign objective (there may be more than one) like selling 20,000 palladium credit cards.
2. **Define the audience –** We then build our segments, and identify targets for the campaign. If we expect a 1 percent response rates, we'd need to identify at least 2,000,000 targets to hit the 20,000 card sales goal. In many cases, we'll use propensity scores as segment criteria, to improve that targeting.
3. **Design the creative –** When we've got our audience, we design and build out the campaign creative. In this case, the creative would be designed to sell the palladium card, and be attractive to the target audience we segmented.
4. **Optimize across campaigns –** If targets are in multiple audiences, we remove them from certain campaigns, and leave them in others – based on constraints regarding number of contacts, required campaign volumes, etc.
5. **Execute the outreach –** We push the campaign out through outbound channels like email, SMS, direct mail, digital, etc, and in many cases set up the offer to display on specific inbound channels, too.
6. **Analyze the response –** Our campaign analysts collect all the promotion and transaction data for the campaign period, and analyzes the attributes of both responders and non-responders, to improve our segmentation, creative, etc.

There are ways in which a campaign can be optimized, better staged, or otherwise improved, but this is pretty much textbook. Depressingly, given the effort, a 1 percent response rate for a credit card offer would likely be *above* average. What's distressing is that companies may well be running hundreds (or even thousands) of similar initiatives – for other credit cards, saving products, mortgages, wealth management, and more, plus a large number of service or compliance messages. These happen in silos, isolated from each other, all trying to achieve their defined goals without regards for other campaigns. This type of thinking is inside-out and product-centric, not outside-in and customer-centric.

The difference between traditional campaigns and next-best-offer campaigns

We've spent a lot of time talking about why campaigns and segmentation don't work very well – but that doesn't have any practical value unless we can recommend a better alternative, and help you understand why it's better. So we're going to introduce the one-to-one approach, which isn't well understood.

You hear a lot about "one-to-one campaigns" or "next-best-offer (NBO) campaigns," which is very misleading and an attempt by campaign management vendors to remain relevant. There's no such thing as a one-to-one campaign – it's an oxymoron.

Let's compare and contrast traditional campaigns with the one-to-one approach:

Traditional campaigns	Next-Best-Offer campaigns
Select the Offer to Push	Select the Customer to Engage
Define the Audience with Segments	Select the Best Offer with Decisioning AI
Schedule Outreach in Outbound Channels	Always On, Across Inbound and Outbound
Campaign Manually Starts and Ends	Offers Activated / Deactivated, Automatically
Success Determined by Offer Response	Success Determined by Change In Customer Value

While campaign management vendors claim otherwise, campaigns are fundamentally different from always-on, one-to-one with NBO. Campaigns focus on the product you want to sell, as well as on your timing. That's why the response rates and experience are dreadful – it's about you, not about the customer. Next-best-offer (or more generally, next-best-action, when it's not just about selling) starts with the customer, first, then uses advanced analytics to engage him or her only when there is something relevant to say, in the right context… when the customer has a specific need, and is ready to listen. One approach is truly the opposite of the other.

Next-best-action (NBA) is an approach that targets individual customers, rather than segments – leveraging their unique needs, preferences, and context. It works to make every interaction relevant and meaningful (regardless of channel), and optimize high-level KPI like revenue, profit, and customer lifetime value, rather than short-term metrics like campaign response rate.

Think about it. How did we get to a place where outbound campaign management teams are disconnected from the targeting strategy used in on-line or digital channels (let alone paid channels), and vice versa? How does that make sense, given those functions share the same customers, and fill in very different, very valuable aspects of the same customer profiles? This situation forces the left hand to act without knowledge of the right, building completely different strategies – a situation which compromises our big-picture goals – because those various strategies are optimized in the interests of the channel, not the overall relationship. No leader would build this scenario out purposely – but it's so common. Indeed, it's easy to see why response rates are spiraling toward zero.

NBA goes beyond traditional marketing, because it looks at all your potential actions, not just the best sales offers for each customer – and there's a giant difference. With NBA, you can absolutely find the best offer for each person, but in many cases selling isn't your best option, at that specific moment, in that specific channel. Instead, you'd be better off taking non-sales actions – actions that better align with that individual's situation (context). For example, introducing a retention plan to a customer who's become likely to churn, or pro-actively offering service when someone is found to be struggling. You may frequently want to negotiate a discount, attempt to collect a debt, or simply thank someone for being a five year customer. We do these things because when you look at the long-term relationship, in specific moments there can be better options than trying to sell. In fact, in many cases your best option might be to do *nothing* – because there's too much risk tied up in that customer, and you'll be selling or trying to retain bad business. Silence can be an undervalued customer treatment.

How does next-best-offer (NBO) really work?

For now, let's keep it simple, and focus on what's traditionally been the top concern for marketers – the ability to determine the NBO. When you're doing one-to-one marketing, instead of just taking a campaign brief and defining an audience to meet a September launch date, you're going to build a strategy that places the customer at the very center. The question isn't: "Should Mary be part of the campaign for this offer?" That's a very product-centric way of engaging customers, and it's inside-out – it focuses on what we want to sell.

Instead, the market demands that we take a more customer-centric approach. So, the question instead becomes, "what offer should we talk to Mary about if we were to engage right *now*?" With NBA, your strategy will first make the decision about which offer is best for Mary in each potential category – which will generally look like this example, for credit cards.

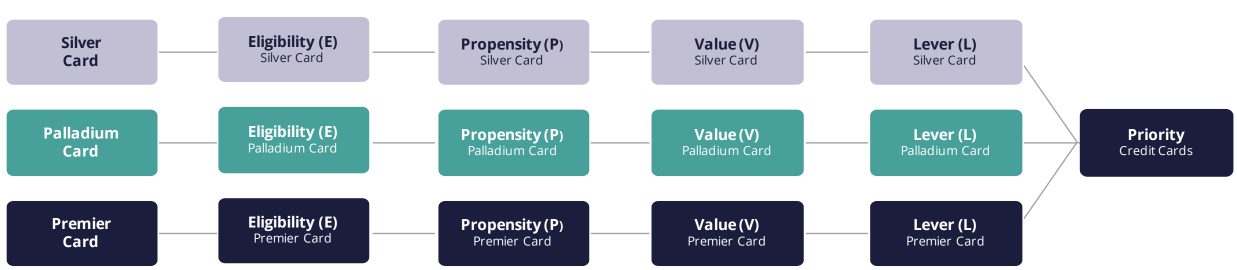

This strategy flows from left to right. For every product, a predictive model will determine the **propensity (P)** for that unique customer to accept a specific offer[4]. **Value (V)** is determined by likely CLV[5] or profit contribution, and rules are used to determine **eligibility (E),** or we can extend the strategy to account for additional factors like suitability, timing and urgency. Finally, there's a **priority** component that selects the best offer from amongst all the options being considered.

When we expand our scope to include all of our potential categories, we repeat this process for each – then compare the winners to find a true NBO, as shown here:

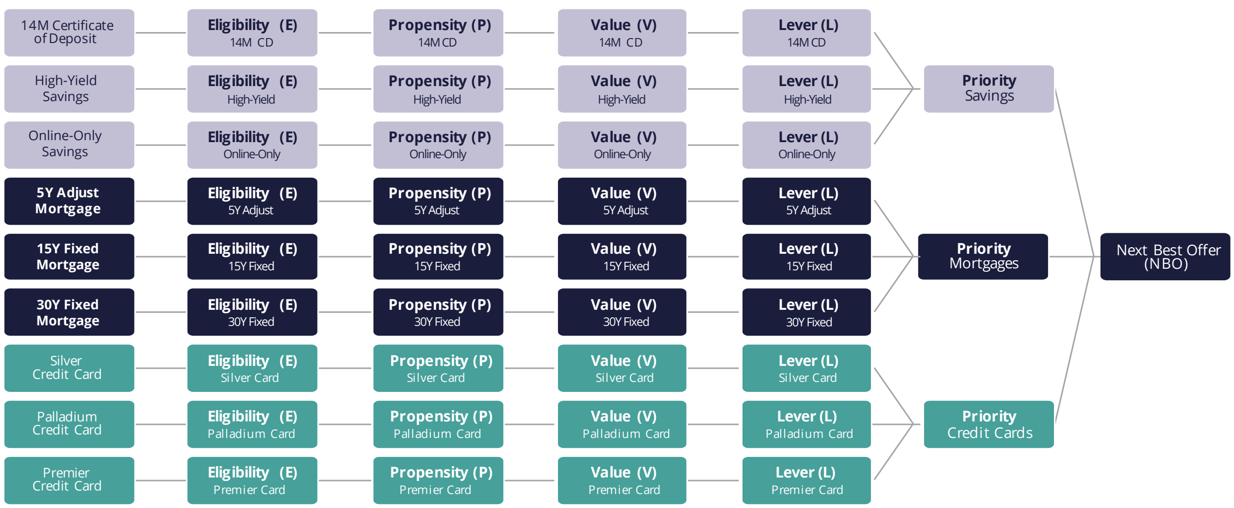

Every company defines 'best' in its own way ¬– but typically it's the offer that will increase customer lifetime value (CLV) by the largest amount, using a formula that, in its simplest form, will look something like Priority = P*V*L, where:

- **P** is the **propensity** of the customer to accept this specific offer. Propensity is a number between zero and one, and is the biggest indicator that the offer will be relevant to the customer, *(e.g. if P = 0.63, there's a 63% chance of that person accepting your card offer).*
- **V** is the **value** to the company when the specific offer is accepted by that customer – it could be a simple margin calculation, a dollar profit level, or a CLV calculation. *(e.g. V = $1,576 in profit when that person accepts a card offer).* Usually, this formula is where a lot of differentiation is achieved, because it's simplistic to equate value with product margins. Doing that would favor high-margin products too much, forgoing 'seeding' offers or messages that would set us up better for a future high-margin transaction.
- **L** is a **lever** used to boost a low-propensity offer, if it makes sense for the business. This is sometimes necessary, but it's tricky – going against propensity will decrease relevance, and will undermine the experience. *(e.g. If L = 1.5, that would boost the priority score for the offer by 50%. An L=1.2 would boost it by 20%, etc...).*

4 If you're thinking this means a predictive model will be used for *every single* offer, you're 100% correct. If you think that's impossible, you're *dead wrong*. For a while now, machine learning (AI) has been able to automatically create adaptive or self-learning models to support / complement data scientists, and scale to meet modelling demand. It's eminently possible to calculate 1,000's of concurrent propensities – and do it efficiently, in real-time.

5 Customer Lifetime Value (also: CLTV). A calculation (or prediction) of the net profit attributed to the entire future relationship with a customer.

When we execute this strategy, the priority rules on the left are applied first – the strategy prioritizes each offer within a product family (the best card, savings, and mortgage offers), then selects the best overall offer using a second prioritization. Because we built eligibility and suitability rules into the strategy, the "best" offer is guaranteed to be suitable, and the customer is definitely eligible for it.

The following chart shows how this calculation breaks down, and how the offers are prioritized by score, to determine the best offer in a category:

Potential Offers	P Propensity	V Value	L Lever	Priority Score
Offer: Palladium Card	63%	$894	-	563.22
Offer: Premier Card	81%	$507	+10%	451.74
Offer: Silver Card	23%	$409	-10%	103.48

Among the three offers that could be presented to this customer, the Palladium Card offer wins, with a score of 563.22. What we've done, ultimately, is pick the offer that best balances what's relevant for that specific customer (using the propensity, P), with what is financially best for the business (using a projected value V for that action), while allowing the brand to boost the priority of any offers it feels deserve a higher priority than the data might show (using the L lever). Again, determining the formula for that balance is an important part of a company's strategy and brand.

Adding new offers to achieve scale

One of the advantages of the NBA approach is scalability – if you need to add a new offer (which is very common), the process to do so is simple, compared to other frameworks. You simply:

1. Define eligibility, suitability, and more criterias for the new offer. This is likely copied from a similar offer, then customized.
2. Define the value (V) and lever (L) priority values for the new offer. This, similarly, is usually taken from a corporate template because artificially high values for V and L will run against our customer-centricity objective.
3. Allow the system to calculate propensity (P) using its adaptive models, or by leveraging existing predictive models built by data scientists.
4. Add the new card offer to the existing strategy.
5. Simulate the modified strategy to see how the change impacts KPI over a period of time – things like expected offer frequency, likely accepts, revenue performance, etc. The simulation shows these metrics in the context of all other active offers, giving you a holistic view of the business impact.
6. If the results are acceptable, the offer is activated. Immediately, it will be offered, when appropriate and relevant, in any channel connected to the hub that's running this strategy (which should ideally be *all* channels).
7. If not, reconsider step two to fine-tune (assuming no errors in step one).

In a strategy, this would mean adding a new offer – labeled as new credit card, below – into an existing strategy for a category, where it would then be considered among all the potential options:

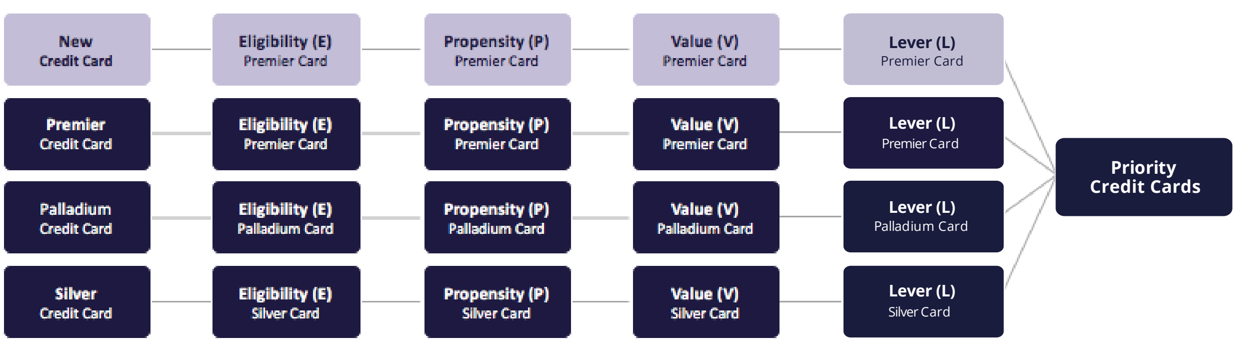

Although this is the one-to-one equivalent of creating a new campaign, the impact on corporate revenues and customer satisfaction is much more profound. This is because, unlike a campaign for a specific offer (a push campaign), in this case the offer is considered along with *all* the other potential options, and we can gauge the opportunity cost of selling it, versus selling other products. And that's just the traditional NBO/marketing capability. In a complete NBA strategy, we could quickly gauge the impact of an offer across *all* the organization's KPI (retention, service, risk, and collection). Well above and beyond the purview of marketing, it becomes a comprehensive customer engagement strategy.

Here's a practical example, courtesy Karen Larrimer, head of Retail Bank and chief customer officer at PNC Bank. The following graphic showcases how NBA strategies add intelligence to their customer interactions beyond marketing, by ensuring that each action taken finds a middle ground that balances their primary concerns of relevance, feasibility, value, and reliability:

Leveraging a single customer decision authority

NBA strategies are designed to materially increase overall customer value. That's what they do, unless we get in their way. But to make that happen across the business, we need to leverage our NBA strategy across *all channels* – both inbound and outbound. Here's why that's important:

- **Unconnected channel conflicts –** Any unconnected channel is a missed opportunity, not only to add intelligence, but also to learn, and to create value. And its activity will actively work against your overall goals. Because it isn't linked to a single decision authority, it will make decisions that collide with the rest of your strategy – pushing offers at the wrong time, ignoring customer context, creating blind spots, and disjointing the overall experience.
- **No optimization of corporate-level KPI –** If several independent marketing/customer engagement strategies touch the same customers,[6] it becomes impossible to simulate their combined impact on the overall company. You can attempt to analyze this after the fact, but only with expensive data aggregation. You're in reactive mode by default. Any organization without active, centralized customer intelligence is running their engagement half-blind – or at best, sub-optimally.
- **Change management nightmares –** A centralized strategy gives you one place to design, test, deploy, monitor, and rollback strategy changes. Having all customer interactions governed by a single decision authority makes change management safer, better informed, and more operationally efficient. How can an organization possibly manage multiple decision-making systems that influence each other, but are managed and executed separately?

While we don't recommend launching in all channels simultaneously (there are significant benefits to an agile, channel-by-channel approach), to be truly effective, the initiatives should expand beyond the real-time inbound channels – through your outbound, event-triggers, paid media, and so on. We'll detail the impact of one-to-one on business outcomes later, but just for context: you can expect a lift in offer acceptance rates of anywhere from 4x, to in excess of 12x more. Even ignoring other benefits, a 4x or higher revenue multiple gives you a compelling business case when you need to connect another channel.

Once you've established a decision authority that crosses multiple channels (a.k.a. a customer decision hub), those channels *immediately* begin learning together, and leveraging the cross-channel data and intelligence to adapt with the customer. The NBO you've calculated for a customer might be right for the current context, but what happens when that situation changes? Regardless of where it happens, that rejection is critical – and it will shift your approach. Presenting offers and learning from the response is a type of context change. Just like in a real-life conversation, the NBO (or NBA) will be calculated immediately after the customer response comes in. In interactive channels, your customer decision hub will be pinged for the NBO multiple times during a single real-time interaction.[7]

[6] They may touch *different* customers. Different marketing solutions may be deployed for different mutually exclusive segments (not in the campaign sense).

[7] Of course, this puts a heavy demand on both the throughput and latency requirements for decisions made by the customer decision hub, but state of the art technology crossed this threshold a while back.

The complexity of NBO strategies: It's all relative

From the outside in, NBA may sometimes look like an intimidating strategy. Yes, it may actually be pretty complex – because your business is complex - but never forget that it's intended to eventually replace all the different "brains", in all the disconnected systems through which you now engage customers. NBA is the lynchpin of your program – connecting and adding real-time intelligence to all your touchpoints, and allowing them to work together when interacting with a customer. In a very real way, it ensures that the whole is greater than the sum of its parts.

Without an NBA approach, this simply isn't possible – think of the difficulty in learning, interacting, orchestrating, and adapting across channels, without a centralized approach. The integration efforts alone would increase complexity by an order of magnitude, or more[8], without the tangible benefits of a unified solution. In practice, using a single customer strategy is simpler, more transparent, and far more efficient than managing disparate marketing applications. If the underlying strategy looks complicated, it is because the business is complicated. There's nothing trivial about balancing customer vs business needs in real time, and optimizing streams of long-term customer value. But in reality, NBA is much simpler, and easier to understand or explain, than the alternatives.

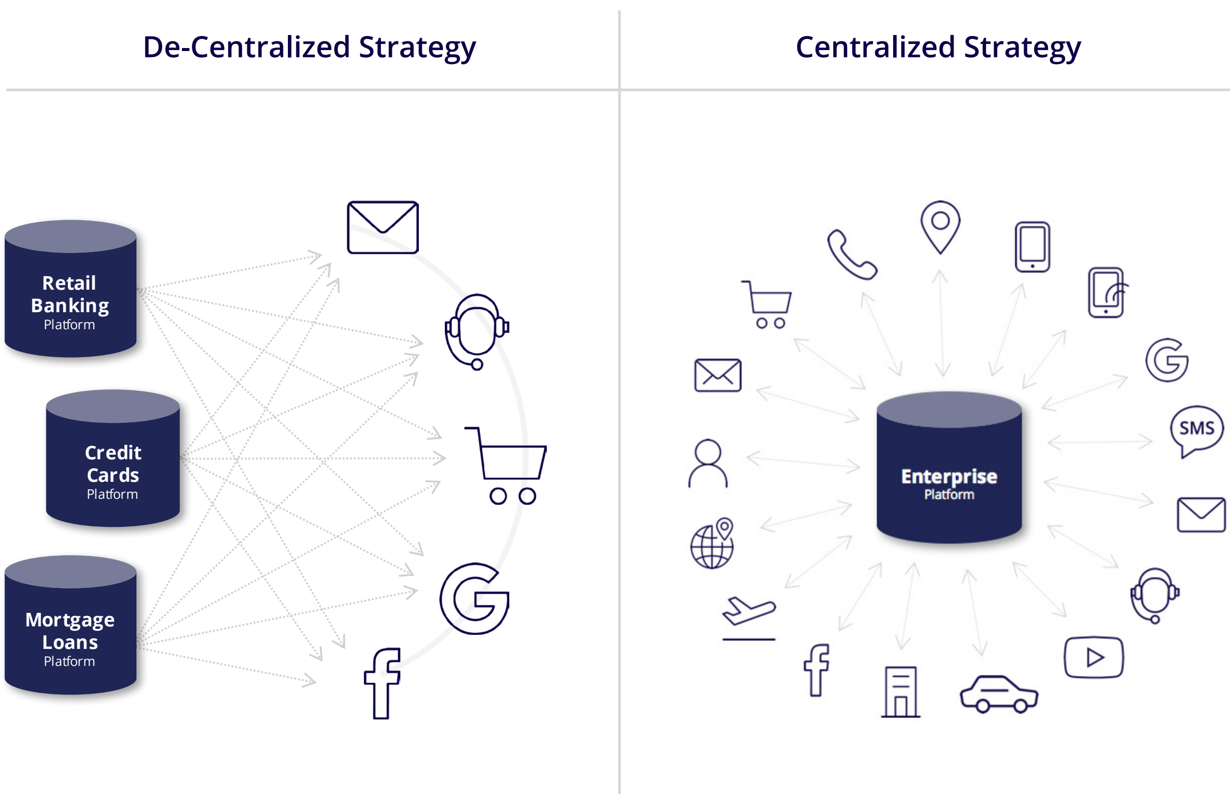

[8] Think of of repeated eligibility criteria, exclusions, and many other logic patterns.

From push-campaigns to always-on engagement

We talked before about relevance, but timing has also become critically important to marketing – and traditional campaigns can't meet the timing needs of a connected customer. If you set relevance aside, the biggest factor in poor response is bad timing – when you run campaigns, you're pushing those messages out on your timeframe, not the customer's timeframe. With all the planning cycles and lag involved in executing a large campaign, there's only a small chance that you'll hit the right window of opportunity to capture an individual customer's attention – let's say a 1 percent chance – so the other 99 percent of what you pushed out gets ignored. You can't wait days, weeks, or even months to fill a customer's needs, once they're presented – consumers are driven by instant gratification, and the opportunity disappears in hours or minutes, rather than days or weeks.

One of the exceptions to this rule are "re-targeting" campaigns, which work because they're focused on timing. For example, a consumer came to your website and viewed a product page, so you dropped a cookie, and fired a pixel – now you can re-market that product to them wherever they go online, via paid channels. Their website visit provided insight into their timing and context – it increased the likelihood that they're in-market to purchase, so it's going to be more effective than a traditional campaign. Now, re-targeting is often seriously over-leveraged, and often executed without precision… but it's a solid example of why the always-on approach works so well – because it's triggered by relevant behavior.

Think about it. Aside from all the effort involved in creating a complex campaign (and that's a significant amount of hustle and bustle), why are you batching everything up, and waiting for specific time slots to launch all that activity? We all know how small our windows of opportunity are… wouldn't it be smarter and more profitable to align what we're saying, with what the customer actually needs in that moment? Wouldn't it provide a better experience if what we said was actually relevant for them, right now? Or if we avoided talking when they're weren't ready to listen? Instead of investing 90 percent of our time building megalithic campaigns to expose a new offer ("look at what I want to sell!"), wouldn't it be far more effective to simply activate that offer, and let the AI/ decision engine figure out when it's relevant to present, for any customer and context, across any channel?

In a traditional model, all of a brand's campaigns essentially "compete" with each other, vying for customer attention – and when they happen within too small a window, each interaction just steals mind-share from the next. The fact that a customer saw one message, or two, or three, doesn't enhance their experience; it actually decreases the likelihood that they'll engage with the next one, especially if the last provided no value. So campaigns are always assigned specific timeslots to minimize customer friction, making it hard to engage in an agile way. This gets compounded by the long, elaborate process required to move a campaign from brief to execution.

Let's really consider this, just for a minute – we know that our campaigns are so poorly accepted, that they're going to interrupt, distract, and maybe even irritate our customers, to the point where we're required to stop talking for periods of time. Is that the best we can really do? Is that how the modern marketer is going to define success?

How always-on marketing is different

Using an always-on marketing approach, none of these requirements or limitations make sense – they don't match the needs of a customer. When we have something new to talk about, the offer or message simply gets activated, and all timing-related rules (contact frequency, customer preferences, channel constraints) become part of the NBA strategy. The offer is "turned on" when it's ready – and presented when it's relevant, wherever the customer wants to interact – across both inbound and outbound initiatives.

If it's relevant and profitable to reach out to a customer about an offer, a one-to-one outbound message doesn't have to wait for a campaign to be executed in – it's simply sent to that person when it's appropriate, and it's the best action to take at that moment. Maybe it gets triggered because the customer's purchase propensity jumped over a threshold, based on what she just viewed on the website. Or maybe she clicked through a Google search ad for a credit card, and we need to follow up. Or maybe she demonstrated a product interest on an owned channel, and we want to re-engage her on a paid channel like Google or Facebook. NBA strategies constantly read the environment, re-assess, and trigger new types of engagement – they don't ignore new information, or context clues from the customer. When we collect new intelligence, we might shift the channel, shift the timing, shift the offer, or just do nothing at all. Ultimately, it depends on what's best for the customer, and the business, in that moment.

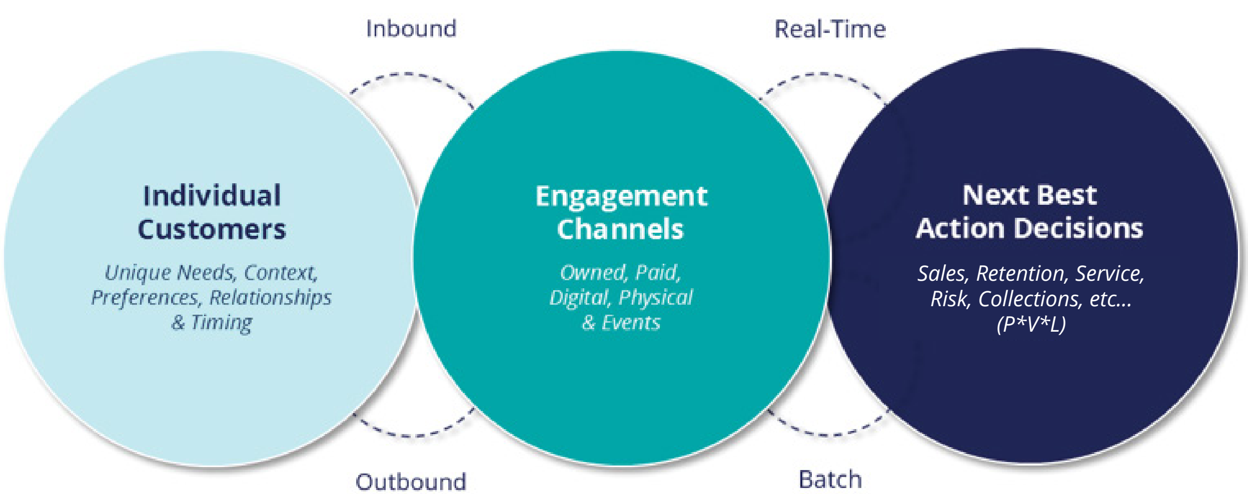

This is the foundation for always-on-marketing – where every customer interaction is relevant, timely, contextual, and in everyone's best interest. This is a very dynamic approach where all interactions (initiated by the customer, or by the company) are used to deliver the right message, at the right time, in the right channel. This 'triple right' phrase is very popular, and highly desirable – but it's only feasible if all three of the core innovations are in place, and working together.

Always-on marketing trigger examples

There are opportunities to communicate with customers every day, and provide value – if you've got the data, analytics, and AI in place to identify them when they happen, and you're agile enough to take advantage quickly. Below are examples of several high-value, trigger-based marketing scenarios, which are being used by companies to identify and engage with customers during their "moments of need" – and not always in a sales context.

These triggers are just examples – in a real-world scenario, AI will actually define the majority of triggers using predictive techniques (and enable us to action them through our customer strategies), so we won't need to make assumptions about which are actually relevant.

Trigger / Activity	Examples	Used To Trigger:
Propensity to Purchase	(P↑ by 30%) & (V >= $250)	Cross-Sell Message
Churn Likelihood	(P↑ by 30%) & (V >= $5000)	Retention Message
Interaction Lapse	No Interactions >= 30 Days	Next-Best-Action Message
Crossing Geo-Fence	Mobil Enters Trade Area	Contextual Message / Alert
Abandoned Cart	(Open > 1 Hour) & (No Purchase)	Re-Targeting Message
Product Expiration	Date > Product Expiration Date	1-Click Replenishment Offer
Upgrade Available	Eligibility = Y, CLV > $2500	1-Click Upgrade Offer
Life-Stage Event	Marriage, Birth, Home, etc.	Contextual Message
Payment Missed	(Late > 30 Days) & (CLV <= $1000)	Reminder Message, Account Lock
Repeated Outage	Outage Count >=3	Service Message, Discount Offer

These triggers will always be 'triaged' by the overall NBA strategy to make sure they further the overall business objectives, exactly like any other consideration to reach out to a customer. For instance, a churn trigger may not lead to a retention effort if the customer is in collections or otherwise perceived as not contributing value.

Optimization of the enterprise - not the unit, product, or offer

A potential downside to the always-on marketing approach is that while it accelerates overall value for the company, results will be less predictable at the business unit, product, and offer level. Think about it - you're no longer pushing products to a set number of customers anymore. Rather, you're finding relevant products, for the customers it makes sense to interact with. Indeed, traditional ways of evaluating "campaign success" are no longer valid. It's not about whether you've met the campaign response numbers by pushing specific smartphones, credit cards, or policies – it's about optimizing the mix of what you talk about, given *all* the potential options. In a framework where you can be extremely precise with relevance, timing, and value projection – it's about how you *balance* what to talk about, and when, to optimize revenue and profit.

Given that, simulation and monitoring are critical to an always-on, one-to-one marketing approach. Before an offer is even activated, you can run simulations that project its likely performance over time. Given all the possible conversations you could have with a customer, it's harder to project how many times one specific offer may be presented – but you can project variations in that number over time, tweak its priority, and see how it performs at the expense of other offers. The end-goal is to introduce new offers into the strategy in a balanced, low-risk way that optimizes value and ensures that when you change a strategy, there isn't a cascade of negative impacts.

Of course, while it's best-practice to let the AI optimize the strategy using propensity, value, and timing, you can adjust the lever (L) in P*V*L to *manually* shift the priority, if it makes sense.[9] You could even have the NBA strategy do this automatically, modifying itself to favor an important offer, when "reality runs behind plan".

But ultimately, NBA will not give business-unit, product, or offer stakeholders the exact same level of control that a campaign would offer them – but that's a good thing. NBA optimization will be better for the company overall, and while a campaign might consistently get you a 1percent response rate, that's all you'll get: 1 percent. That's a massive compromise, given the response and performance potential available.

[9] It often doesn't make sense, but there can be good business reasons to play with 'override' controls like the lever L. E.g. to counter a competitor's actions, or for a brand awareness rationale that's not a simple thing to monetize objectively.

Real-world value: The proof is in the pudding

Anyone can chatter on about a concept like one-to-one marketing. Unless, however, the underlying theory can be converted to actual business value, and provides a superior alternative to the status quo, then it's simply a distraction. Companies like the following have made the leap, and moved from theory to practical reality – with transformative impact:

RBS
The Royal Bank of Scotland

Before the 2008 recession, **The Royal Bank of Scotland** was the largest bank in the world – but also the least trusted bank, in what was the least trusted industry in all of the U.K. They invested in their "Always On Customer Brain" to reconnect with customers. They now message customers 35 percent less often – but have generated a 600% increase in response rates, and produced a 10-1 overall ROI.

Elisa is a 130 year-old telecommunications company and the wireless market leader in in Finland, providing a wide variety of online services internationally. A Pega Next-Best-Action customer since 2009, they've progressively integrated their inbound & outbound channels, demonstrating a 20x increase in sales opportunities, 150x increase in banner engagement, 5x increase in telesales conversion, and a 6x increase in email click-through rates.

CommonwealthBank

The largest bank in Australia, **Commonwealth Bank** provides best-in-class experiences to 13.5 million customers through what it calls "Next Best Conversations". They constantly re-assess each customer's context to power 5 billion interactions per year, across both inbound and outbound channels. They've seen a 10x increase in lead volume, and a 3x increase in conversion.

OPTUS yes

Optus is Australia's second largest telecommunications provider, with greater than 10 million customers in a country of 22 million people. Needing to get faster and smarter than competitors (but without data science resources), they implemented in the NBA framework for their outbound marketing channels, with real-time decisions – reducing campaign time-to-market by more than 90 percent, from five weeks to less than 24 hours.

British Gas

British Gas is an iconic U.K. brand with an over 200 year history – the first public utility established worldwide in 1812. They needed to move from a "product push" approach, to a data-informed "customer conversation" – with one system that could handle both inbound and outbound. With all of their channels now connected, they simulate the performance of 1billion customer offers in less than three hours.

From next-best-offer to next-best-action

Advancing from an NBO approach (for marketing) to an NBA approach (moving into customer experience), is systematically very straightforward – but requires the logic to extend to take into account higher-level considerations beyond sales – such as retention, service, risk, collections, negotiation, and so on.

Reading the strategy below from right to left, we see there is an explicit definition of the company's business objectives, and trade-offs – and like with NBO, a prioritization that decides when each of those business strategies are relevant.

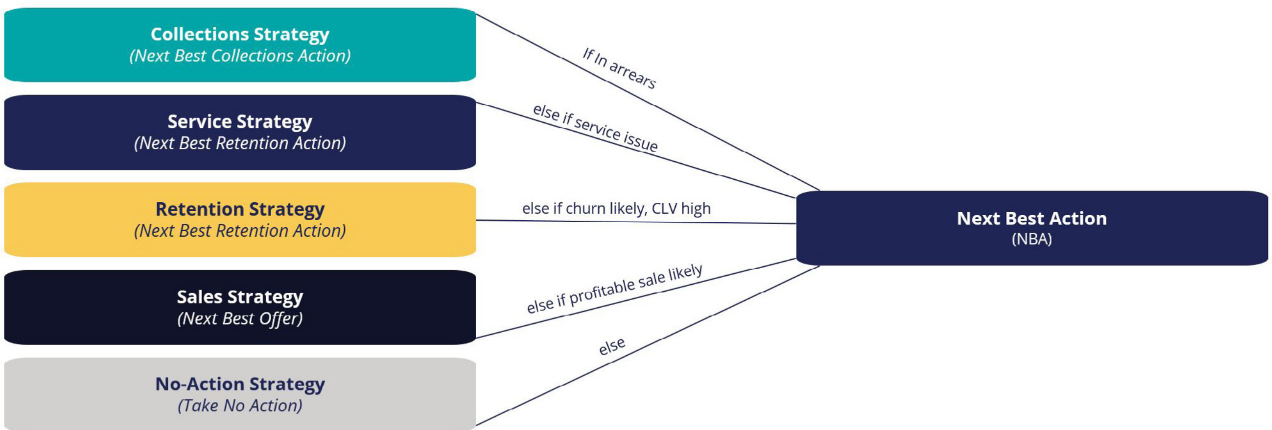

For instance, instead of simply considering retention plans against sales offers in every scenario, Next-Best-Action will shift to retention only when specific conditions are met – such as a customer's "likelihood to churn" score rising above a certain threshold. Other business concerns, such as risk, can recommend that we do *nothing* when they're invoked – because the risk of doing business with a specific customer has become too high.

The roadmap to always-on: Best practice options

Some enterprise-level organizations maintain hundreds of millions of customer relationships, with billions of complex interactions each year – spread across dozens of marketing and engagement channels…nothing about that is remotely simple. So the decision about how to organize and advance the move from segments and campaigns to NBA and always-on requires some strategic thinking. However, it's a great opportunity to transform corporate thinking and culture towards true customer-centricity, funded, every step of the way, by massive business outcomes improvements.

Organizations will attempt to answer the following questions (and assemble a great deal of data and consensus) before making a decision on their always-on roadmap:

- Which channel or business function is causing the most customer pain, currently?
- Which would yield the greatest revenue gains? Cost efficiencies?
- Which would generate the greatest expense? How do we optimize cash-flows?
- Which has the greatest switching costs? Would those decrease over time?
- Which stakeholders would be the most likely to adopt? Where would there be resistance?
- Which channels provide data or insights that can be leveraged downstream?
- How will our organizational structure support this initiative?
- Who will champion this initiative – now, and after the initial implementation?
- How can we generate quick wins to validate the project, and build momentum?
- What order of operations has produced success? Are there established best practices?

Real-world roadmap examples

There are hundreds of possible variations in the roadmap for an NBA implementation, as an organization converts quick wins and then expands across channels – those will vary greatly depending on the nature of an organization's products, industry, levels of expertise with AI/analytics, technological capabilities, go-to-market strategy, and customer orientations.

The following are actual real-world examples that have been repeatedly adopted, templated, and modified by other organizations. Each approach provides value because:

- They provide tangible quick-wins that substantiate the investment early-on.
- They showcase the potential of NBA with a robust path to value.
- They work to mitigate growing pains that could generate undue stress for the core business.

Roadmap 1: Opening with online channels

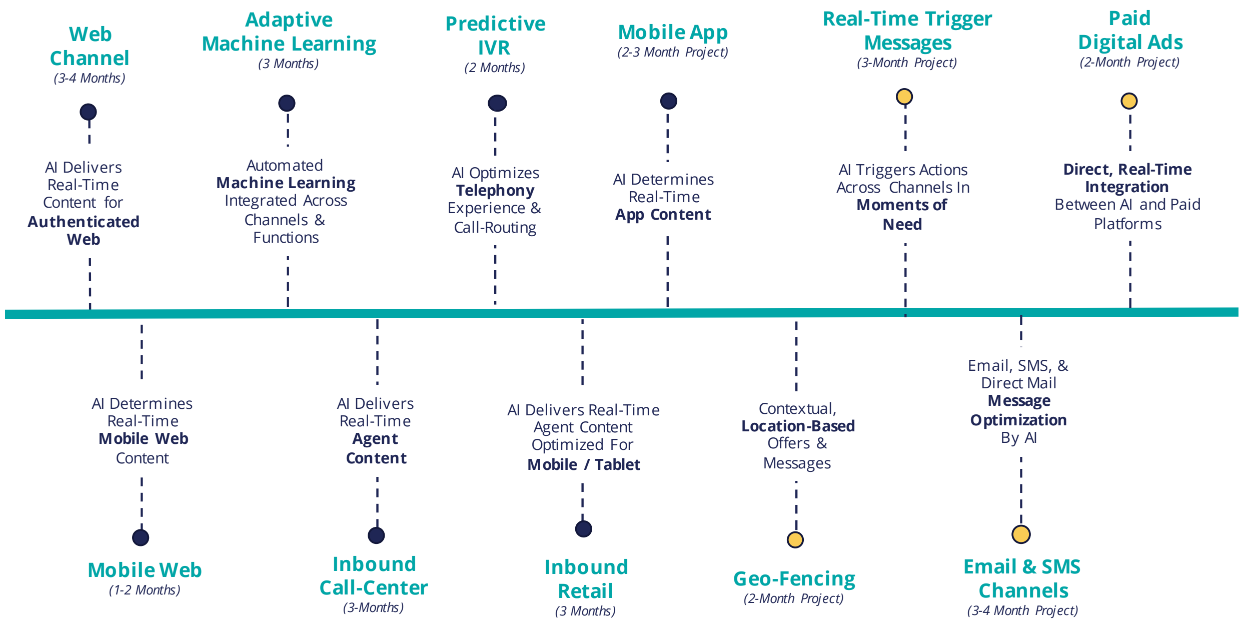

- **Roadmap 1:** Opening With Online –This path prioritizes digital/online channels, specifically the web and mobile.

Roadmap 2: Activating agent channels

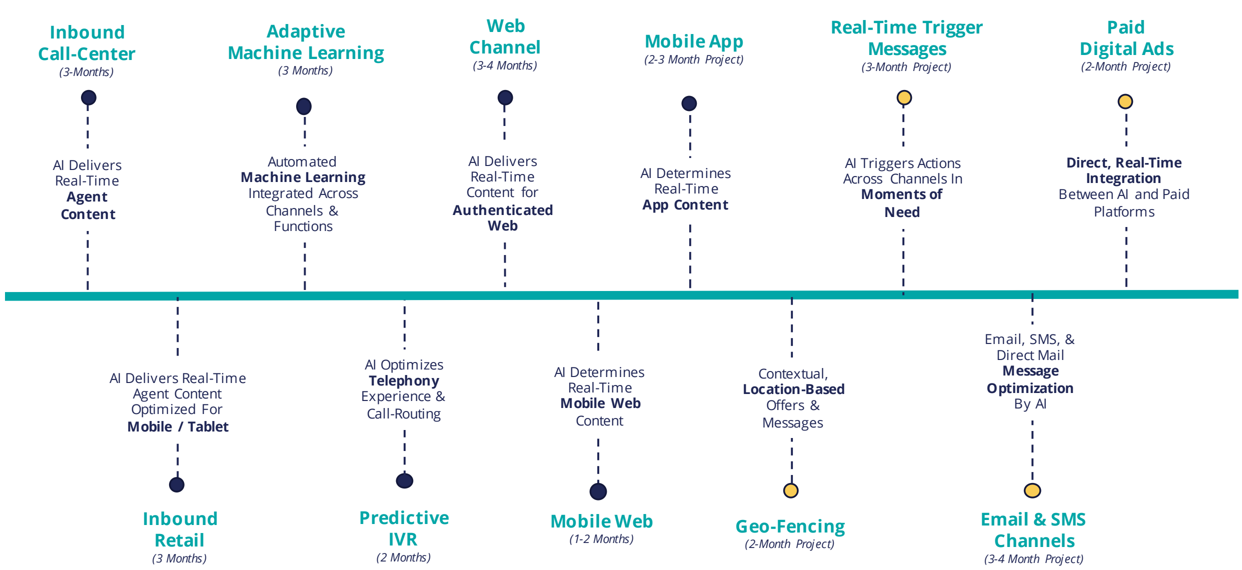

- **Roadmap 2:** Activating Agent Channels – This path focuses on the integration of agent-facing channels, such as the call-center or retail locations.

Roadmap 3: Ramping up with real time

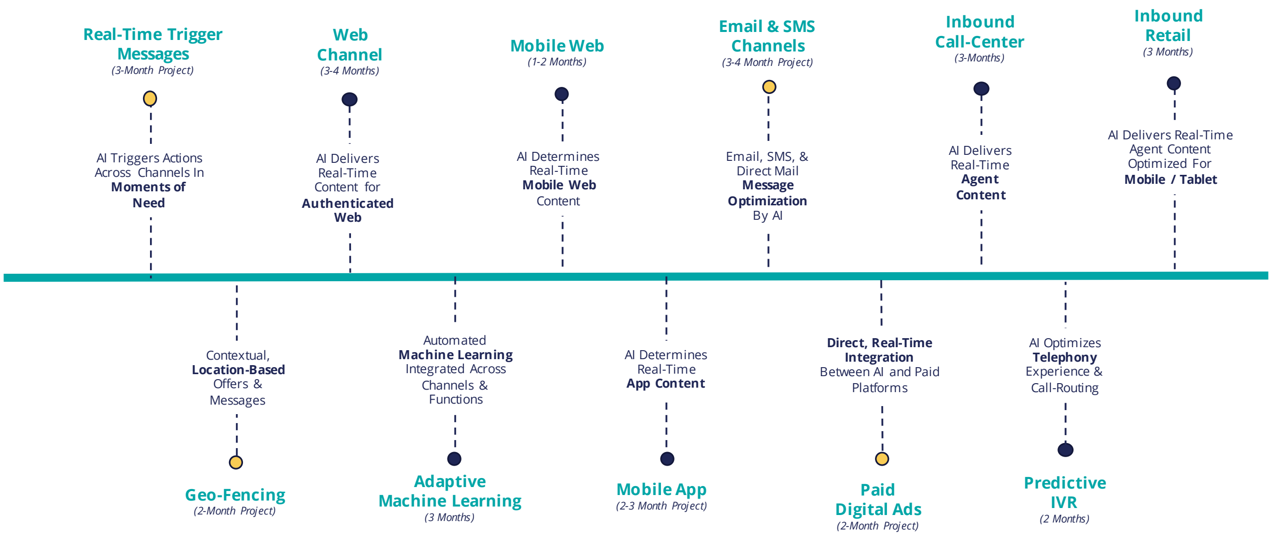

- **Roadmap 3:** Ramping Up with Real-Time – This path champions pro-active always-on triggers, and moves quickly into a mix of inbound and outbound channels.

Conclusion: Massive benefits from one-to-one approach

Overall, the NBA approach has tremendous benefits over a campaign/segment model, and significant benefits even over the sales-centric NBO – because it allows the organization to prioritize actions designed to boost customer experience and satisfaction, rather than just marketing or sales objectives.

Typically, a comprehensive investment in NBA yields benefits like:

- Increasing response rates (3-6x).
- Reducing churn rates (10-50 percent).
- Reducing retention discounts (20-35 percent).
- Finding Incremental agent-sales opportunities (2-3x).
- Increasing net promoter score (10-40 points).
- Maximizing return on investment (3-5x).
- Minimizing payback period (four to six months).

But again, to make those benefits a reality, the following innovations need to happen within the organization – both technically and culturally:

- The organization must progress from one-to-many segmentation, to a one-to-one customer engagement approach.
- The organization must migrate from channel or product-driven targeting, to a centralized decision authority.
- The organization must reduce reliance on scheduled push campaigns, and advance to an always-on model that engages customers during their "moments of need."

These three innovations are eminently possible today and early-adopters have seen massive gains – demonstrating KPI lift across the board while their status-quo peers were losing ground in the same markets. Given the state of marketing in general, it's clear that we have to change how we think about customers, and be willing to adapt – to invest in an approach that will develop relationships, instead of eroding them.

ABOUT PEGASYSTEMS

We are Pegasystems, the leader in software for customer engagement and operational excellence. Our adaptive, cloud-architected software – built on the unified Pega® Platform – empowers people to rapidly deploy, and easily extend and change applications to meet strategic business needs. Over our 30-year history, we've delivered award-winning capabilities in CRM and BPM, powered by advanced artificial intelligence and robotic automation, to help the world's leading brands achieve breakthrough results.

For more information, please visit us at **WWW.PEGA.COM**