Evaluating next-generation planning & scheduling platforms:

8 Questions to Ask

optibus

Introduction

New mobility alternatives, applications and software are changing the world of transit in unprecedented ways.

In the public transit sector, older methods of planning and scheduling are making way for newer platforms. These new platforms carry enormous transformational potential. They improve operational efficiency, lower operational costs and create novel opportunities for newer, data-driven services.

How can transit providers, government or privately-owned, determine which planning and scheduling solution is right for them? What are the features that can change business for the better?

Whether you're looking to plan better routes, timetables, ensure on-time performance, create better vehicle/crew schedules or rosters (rotas), you will want to ask these **eight questions:**

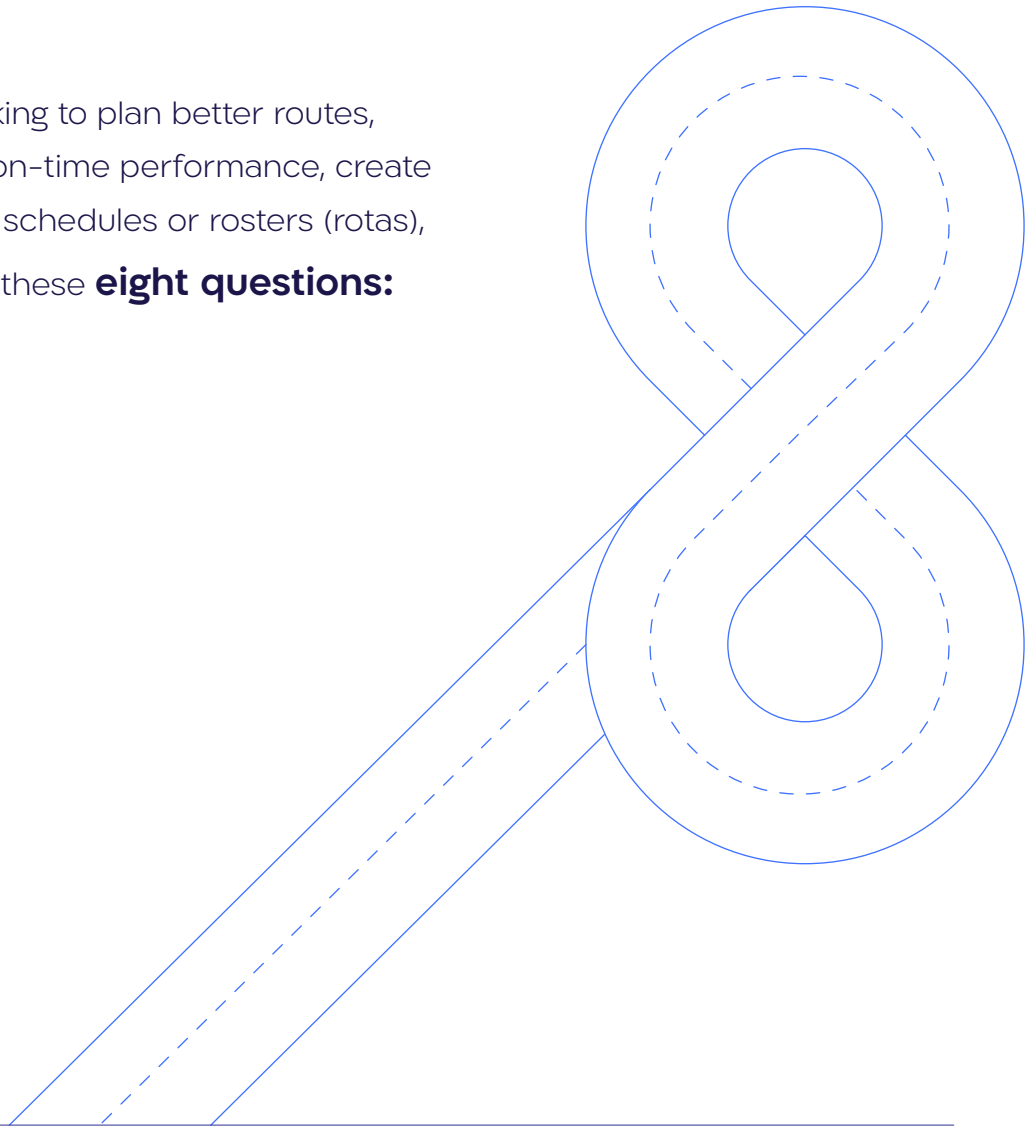

1 Is it easy to accurately express rules and preferences?

Public transit operates under a complex set of rules and constraints: from break rules to depot capacity, from changeovers to interlining etc. The ability to accurately express these requirements is probably the most important factor influencing the success of a planning and scheduling platform.

When rules or preferences aren't precisely expressed before the optimization run, they need to be taken care of later. This is no simple feat: work that needs to be done post-optimization always requires intensive manual editing, from adding relief vehicles to creating rosters. Even worse, this post-optimization work can and will actually degrade any positive business results that optimization was supposed to deliver in the first place.

Cost data for any given scenario is only accurate if it can reflect that scenario accurately and in full. If it can't reflect overtime, relief vehicles, or any other occurrence, it stops being accurate and becomes a rough rule of thumb. When costs aren't precise, they can't be used to compare different optimization scenarios.

That's why it's important to be able to express all rules and constraints, to have a schedule that reflects them in full, and to have the cost data and business KPIs accurately reflect the costs associated with any scheduling scenario.

Validating timetables, routes and schedules is also important. Modern systems use the same mechanism to notify users when schedule or timetable changes are contrary to rules or don't make sense, minimizing the cost of human error.

Does the platform show all relevant business metrics?

Planning and scheduling isn't about checking just one scenario: it is about checking multiple scenarios and picking the best one. Each scenario is different – and each creates different cost, utilization, deadhead and other results.

The best way to determine the optimal scenario is to compare scenarios using relevant and accurate business and operational metrics, knowing they reflect all rules and constraints.

Advanced platforms can also show the impact of changes on business metrics in real time: when manual edits are performed, the platform recommends the ones that have the least impact on the business KPIs, saving scheduler effort and ensuring quality results. Additionally you can add any bespoke business operations metrics that are important to you.

How long does it take to generate a schedule?

Speed is crucial for planning and scheduling. Scheduling is notorious for its lack of speed - it's a complex problem that requires a lot of computing resources.

That's why (for most legacy solutions) optimization runs take hours or days. This is bad news for business agility, since tenders and service changes require a lot of work, and in some cases, schedulers and executives don't have the time to thoroughly check all what-if scenarios and may settle for less than optimal schedules.

That's why next generation planning and scheduling platforms are all about speed, with optimization runs that take minutes or seconds, using parallelization.

These systems were built from scratch so that schedulers and planners can test multiple what-if scenarios and arrive at the best possible scenario, delivering both better passenger service and lowering operational costs.

When thinking of speed, another important factor to consider is how quickly it takes to onboard the system and deploy it - weeks or months?. Afterwards, the next question is how quickly can users begin using the system. This is determined by the modernity and simplicity of the user interface, and the ability to express rules and preferences without any coding or scripting expertise. Ideally, such systems should support an ordinary business user with strong visualizations and use business KPIs as a guide during the optimization process. Ideally, you should look for a system where you can be up and running within a few weeks at most.

4 Can the platform work at scale?

Optimization problems are difficult; they become exponentially more difficult as we grow the number of vehicles and drivers within a given schedule.

Yet scheduling is something that happens at scale: it concerns the public transit of cities with millions of passengers, a high degree of complexity and constant change, across hundreds or thousands of vehicles.

As opposed to legacy platforms, modern systems can easily deal with scale increases - involving sites with hundreds and even thousands of vehicles and drivers. These modern platforms use the cloud (and its ability to parallelize tasks), advanced algorithms and AI to solve massive problems in a short time.

Is the platform data driven?

One of the key assets new mobility providers use to their advantage is data about the movement of people and vehicles. **The underlying data drives innovative services that are based on predicting demand, travel times and more.**

Today public transit providers don't always use the data they already possess. AVL data can be used to predict on-time performance, but it isn't used often - and the interface isn't simple. Data about passenger demand, travel times etc. isn't used as it should.

Next generation planning and scheduling platforms are data driven, taking in multiple sources of data and simply integrating them into the scheduling process. For instance, you can take real-life data, apply artificial intelligence and then predict the likelihood of on-time performance for each element of a schedule.

Combining this with business metrics, the platform can recommend how to reach on-time performance goals at the lowest operational cost.

Last but not least, these platforms should provide robust analytics and granular reporting, that can be customized if needed, to help the users better analyse all the data available to them.

Is the platform web and cloud-based?

Traditionally, legacy planning and scheduling platforms are purchased and installed on the buyer's servers. Since optimization is a heavy operation, buyers often set up server clusters to support more users, and a dedicated database is needed too, generating IT overhead. Even when the software is "cloud-based" (i.e. allowing remote desktop access to servers on the cloud), the resulting optimization runs are slow.

Modern platforms are truly built to be software-as-a-service. This has some important benefits: providing continuous delivery of new versions, features and software upgrades. As a result:

- **Software upgrades and new features are immediately available on the platform - with no need for maintenance or downtime, making it easier for users too.**
- **No IT resources are required for a database, servers etc**
- **Optimization runs are much quicker since they run on distributed cloud assets that are allocated according to need, exceeding the output of large server farms yet requiring none of the maintenance**
- **Collaboration between different teams or management levels is simple and quick**

Does the platform support collaboration?

Collaboration means that different users, in different locations, can work on the same schedule together. This means that the site that creates a driver schedule twice a year can easily get support from a company expert in a different location, and that control and oversight (as well as unified requirements) for different locations are simple to accomplish.

Schedules can be shared with different stakeholders:
agencies, operators, executives etc, and several schedulers can jointly work on the same schedule. Typically this comes with a permission control system, so that the sharing can happen without risk of deletion or mistakes.

Does the platform offer a comprehensive end-to-end solution that is also future proof ?

Mass transit is changing rapidly - it can become demand-responsive, be integrated with other modes of transportation, derive more insights from data about people and their movement in the city and more.

Whatever those changes may be, any next-generation planning and scheduling platform needs to offer a comprehensive end-to-end solution in which all modules offer a seamless workflow and the user can go from route planning through timetables, vehicle and crew scheduling, all the way to rostering, using a sole interface in an all-inclusive solution that offers the flexibility to use existing data sources or start planning a route from scratch. This type of holistic solution will inevitably become the backbone that allows for future developments to be deployed easily and for mass transit providers to take the necessary next steps in improving their operations.

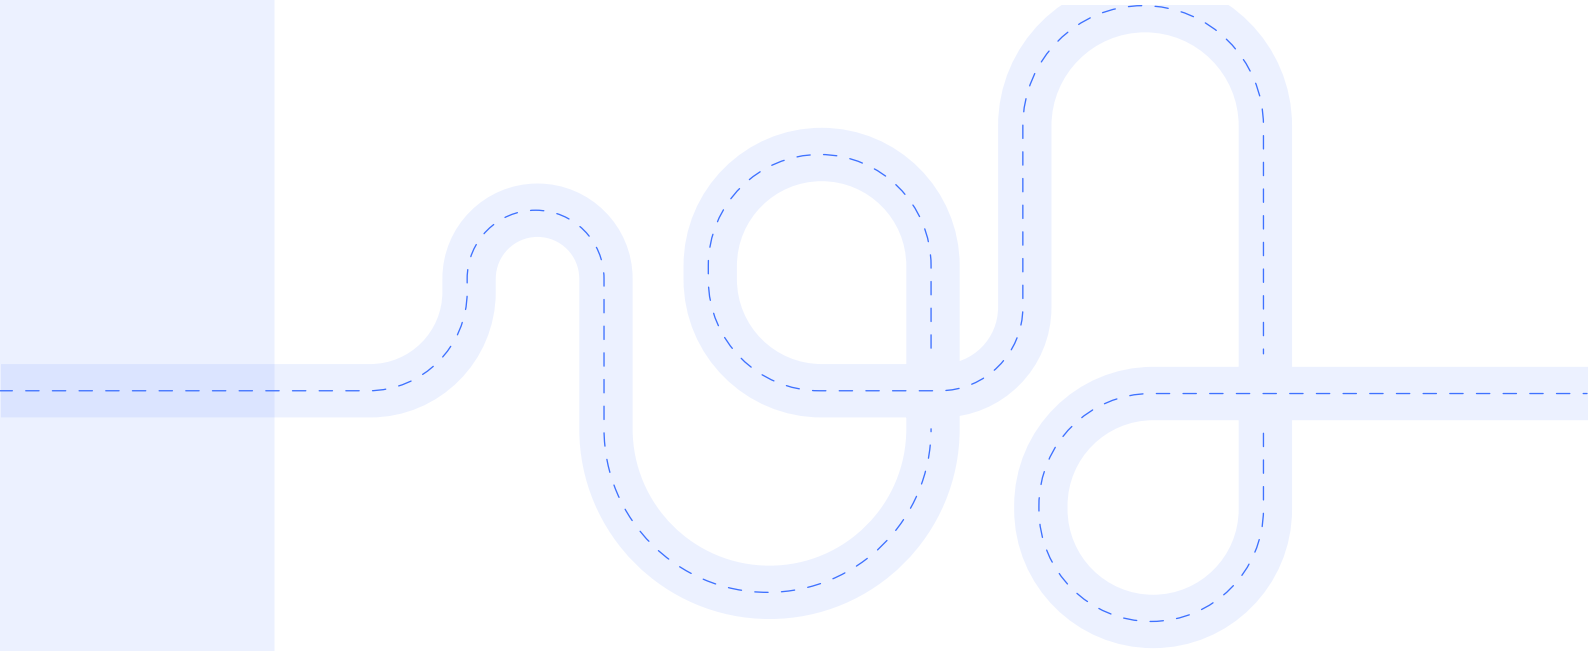

Summary

Mass transit is facing some exciting possibilities, yet they require mass transit operators to undergo digital transformation and enter a world of quick, modern and accurate software-as-a-service platforms.

In the case of mass transit, modernization and digital transformation aren't "nice to have" features: given the enormity of the technological challenge that's inherent in scheduling and planning, these features are a must for any organization that wants to partake in the future of mobility.