Capacity Planning

Predict, forecast, and optimize your workforce in real-time

IFS

Matching team resources with fluctuating demand over multiple time horizons is a tough challenge.
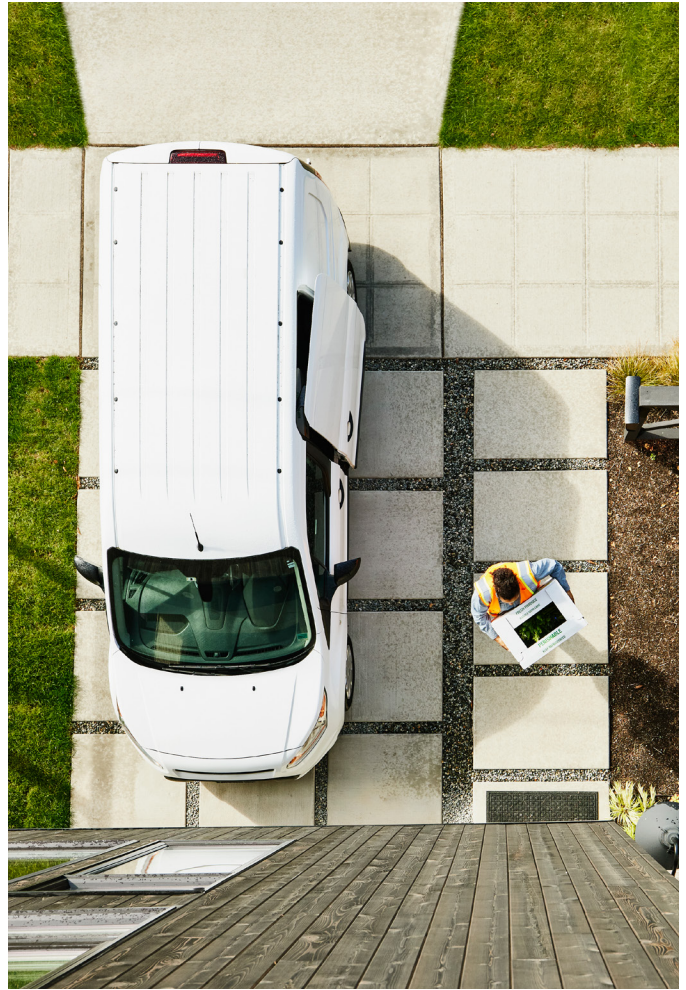
Organizations need real-time flexibility for scenario planning, coupled with forward visibility to strategize effectively and provide best-in-class service.

In this technology spotlight, we examine different mechanisms to optimize capacity planning, including the ability to increase automation using real-time and operational time horizons.

A series of flexible parameters can be used to achieve this result. These parameters absorb incoming workload imbalances by automating capacity and the movement of resources. Less powerful systems use a different model with capacities or buckets set for different work types in a tightly defined regional policy.

However, capacity planning must be flexible, extending visibility to support automation of the schedule weeks and even months into the future—with the necessary tools to adjust a plan based on differing resource levels over tactical and strategic time horizons.

This flexible framework provides greater automation. Schedulers can support demand and resource fluctuations as they happen with much better outcomes for customers.

1	**2**	**3**	**4**
Real-time	**Operational**	**Tactical**	**Strategic**
0–7 days	1–4 weeks	1–6 months	3–12 months
Dynamic scheduling	Appointment booking	Scenario planning	Cyclical maintenance
Always optimized	Shift planning	Resourcing decisions	Resourcing & KPI's
Target exceptions	PM management	Skill requirements	Contingency planning

Flexible work zones

Scheduling engines continuously optimize in real-time based on events happening in the field and elsewhere. This process of ongoing optimization works best with flexible work zones.

However, many batch schedulers follow a common practice that enforces the separation of scheduling problems into small rigid work zones or districts. This enforced separation of data creates inefficiencies and heavier administrative work. It also leads to imbalances with incoming or projected workloads that require manual team changes and/or work reallocation within these tightly defined areas.

Greater flexibility is possible if you soften these boundaries. For example, keeping field technicians "on patch" while also enabling them to work in neighboring zones will automatically counter workflow imbalances.

This approach is less disruptive to the operation while improving the customer experience:

1. Faster response time for urgent, high-priority work

2. More timeslots for appointment booking due to a larger pool of field technicians

3. Minimal back-office administration with less reliance on manual workarounds to cope with fluctuations in demand
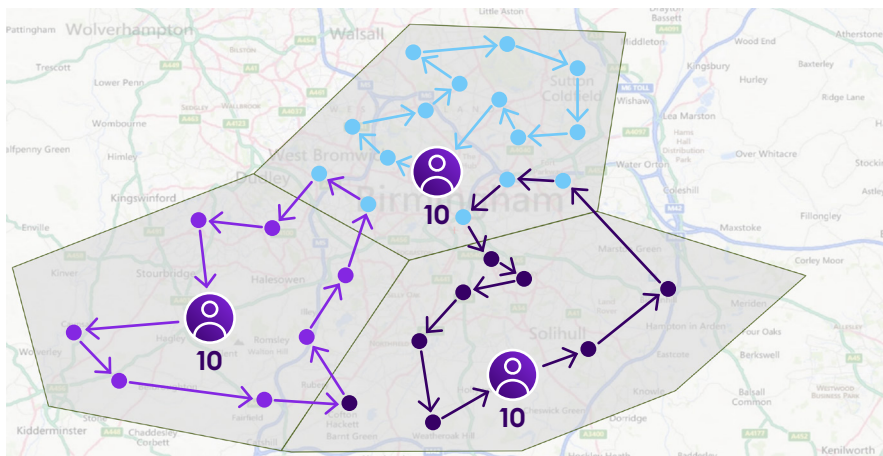
Hard boundaries—no cross over
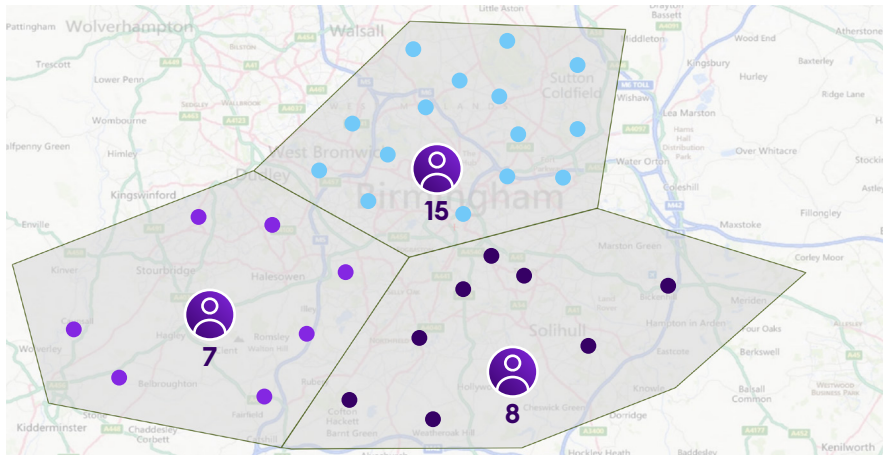
Capacity

10 jobs per day

3 x 10 = 30 jobs

Actual

10 + 8 + 7 = 25 jobs

16% lost productivity

Soft boundaries—flexible, automated

Capacity

10 jobs per day

3 x 10 = 30 jobs

Actual

10 + 10 + 10 = 30 jobs

All jobs complete

Travel cost profile

To maximize productivity, remove regional barriers and policies. Instead, constrain resources with a flexible travel cost model that includes configurable travel time.

In this model, field technicians work within overlapping areas using a travel limit that is acceptable to the business. Criteria are set differently for urban and rural teams, maximizing productivity and operational efficiencies for the organization and your customers.

To achieve this result, your data model must contain parameters to automate flexibility in travel limits. For example, permission to break a travel limit if an SLA is at risk but not for a routine job of less urgency.

Shift utilization

Use a template to set utilization at different levels–for resource groups or down to the individual resource.

The utilization template limits the amount of planned shift time, working on a rolling basis into the future. For example:

- Day 0 100% (today)

- Day +3 80%

- Day +7 60%

- Day +14 onwards 50%

With this level of control, space is held back and made available in a staggered fashion. This model works well for short, reactive SLA assignments that must be combined with longer term planned maintenance, regular visits, or booked appointments.

Appointment booking

Appointments must seamlessly integrate with other work types, while optimizing resource flexibility and efficiency. Best in breed systems include:

- Decision-making that incorporates existing routes and travel times

- Continuous optimization that factors in skills, regions, preferences (preferred/blocked), resource types, costs, etc.

- Available appointment slots based on existing work model

- Slot ranking or scoring based on the cost to perform (pricing logic)

- Slot differentiation based on type of work / customer / contract

- Incorporation of reactive work that aligns with utilization rules on shifts and slots

- Displacement of existing, low priority work as needed

- Premium versus standard services based on job type constraints, including same day appointments

- Overlapping appointment slots

- Dynamic adjustment of existing work allocations to find an optimal slot

- Ongoing optimization of existing work after each slot is booked

These best practices ensure potential appointment times are calculated and ranked based on existing workloads and business policies.

Only appointment slots that accommodate the work will be offered. Considerations include the duration of the job and travel time.

As the schedule fills, slots automatically reduce in real-time. Slot ranking incorporates specific parameters based on how each slot can be accommodated into the schedule.

The appointment slots are defined by a series of templates that consider time, region, and job type.

Appointment templates

Appointment templates determine which timeslots are offered to each customer for different types of work based on a range of criteria you can select:

- Number of slots offered
- Day pattern
- Window length
- Granularity
- Minimum delay (1st offer = now + "n" hours)
- Time of day
- Utilization rules
- Slot usage rules
- Blocking / non-blocking
- Start- and end-based

Appointment templates also align with SLAs for each particular job.

Advanced appointment booking rules

Utilization and capacity rules provide you with precise control over your upcoming workload and resources.

Overbooking
Overbooking is common when a large percentage of the work is appointment-based. When cancellations occur—and with fewer long-term jobs on the schedule to fill the openings—white spaces appear in the schedule.

- A best practice is the provision of a bucket resource to soak up extra work when appointment slots are overbooked:
- Bucket resource holds overbooked jobs
- Preference controls use the bucket resource is a last resort for job allocation

When white space appears, the job is automatically re-allocated to a regular team member

Activity rule collections
Incorporate rules to control any metrics attached to a job, i.e., value of the job, duration, number of jobs per shift, etc.

- Apply rules by job type, defined resources, and/or resource types
- Create a limit based on a configurable time period (shift/day/week)

Job type A

Appointment slots are offered same day and for the next three working days on an AM/PM basis.

Job type B

Appointment slots are offered in 48 hours and may include 2-hour slots into the evening and weekends.

Once a slot is chosen, the SLA for the job adjusts to the appropriate timeframe and the job proceeds within its new constraints. Until the job is completed, it could be re-allocated to a different technician as the system continues to optimize.

Slot usage rules

Apply slot usage rules to appointment slots within an appointment template. This can be done on a rolling basis into the future, similar to the shift utilization rules.

- Offers incorporate the percent allocated into the future

- Appointments will not be offered beyond this limit

- Different settings for different job types

- Premium and standard service levels are applied to different customers

- Capacity is measured across all resources

- Capacity rules are applied across multiple slots

With so many flexible parameters at your disposal, you can precisely tailor offers to your customers to optimize their moment of service. When plans must be adjusted, use tools that help visualize outcomes so you can make preemptive changes.

Changing the plan

Continually review your schedule ahead of time. Determine utilization rates and whether or not capacity must be adjusted.

The scheduling engine continuously optimizes the dynamic window during operational hours. Outside of this time—during non-operational hours—the system uses a full set of algorithms on the larger appointment window to ensure the overall quality of the schedule is maintained.

Filters and zoom controls in the user interface allow you to view the schedule over specific timeframes and areas. Authorized users monitor and escalate exceptions so that action is taken preemptively.

Activities are booked accurately on known capacity. In addition to changing the parameters, increasing or decreasing available capacity will change the availability of appointment slots for customers.

Planning occurs via a dedicated UI where users can create and plan for a variety of shifts many months in advance. Detailed analysis reports are also available to aid in planning.

Add resources, generate and change shifts for individuals and groups, add or modify overtime. When the capacity is changed in any of these manners, extra slots will be opened or diminished accordingly.

"PSO is a phenomenally powerful tool. It is key to us delivering the outcomes our customers want in the most efficient way possible."

Mike Gosling, IT Service Platforms Manager, Cubic Transportation Systems
Read the case study

Strategic forecasting

Capacity planning over the strategic time horizon requires a different approach. **What If Scenario Explorers (WISE)** are critical for this work, allowing you to model scenarios such as shifts in demand, promotional activity, seasonal trends, new contracts, organizational changes, and the removal of regional boundaries.

WISE allows you to predict how changes impact your existing team so you can adapt your field force and maintain optimal customer service levels. The technology helps you answer important questions:

- How many workers are required?

- What skills are needed and where?

- How will a new contract impact service delivery?

- What happens in an extreme weather event?

- What is the impact when work practices change?

- How much travel is required?

- What KPIs are achievable?

- How are field teams adjusted after a merger or takeover?

The scenario explorer uses advanced predictive analytics and provides easily created simulations to help you identify the best path forward. In each scenario, you can interrogate the schedule and examine the forecasted outcome, comparing different scenarios objectively before following a course of action.
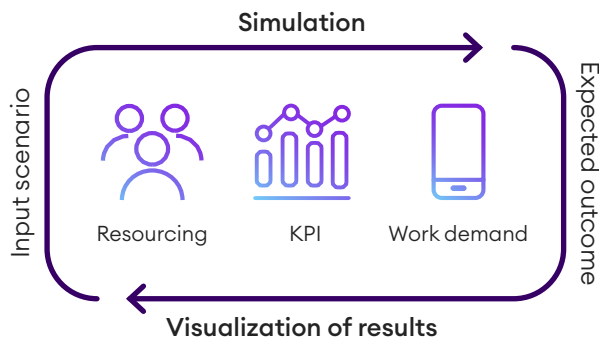
You can use a WISE for assisted and automated planning:

Assisted planning

- Import workload and current staff data

- Forecast KPI performance of current resource

Automatic planning

- Import workloads and current staff data

- Define desired KPI performance

- Forecast resource changes required to meet your target

Summary

Effective capacity planning requires innovative technologies and tools applied across multiple time horizons to efficiently plan and schedule your field teams.

IFS works with organizations worldwide, helping them implement a framework for capacity planning with continuous and automated optimization. The ability to plan scenarios over a longer time horizon allows our customers to make changes preemptively and plan for the future with confidence.

Visit our website to learn more about capacity planning and optimized scheduling or **contact us**.

About IFS

IFS develops and delivers enterprise software for companies around the world who manufacture and distribute goods, build and maintain assets, and manage service-focused operations. Within our single platform, our industry specific products are innately connected to a single data model and use embedded digital innovation so that our customers can be their best when it really matters to their customers—at the Moment of Service™.

The industry expertise of our people and of our growing ecosystem, together with a commitment to deliver value at every single step, has made IFS a recognized leader and the most recommended supplier in our sector. Our team of 4,500 employees every day live our values of agility, trustworthiness and collaboration in how we support our 10,000+ customers.

Learn more about how our enterprise software solutions can help your business today at ifs.com.

#MomentOfService