Workload Partitioning

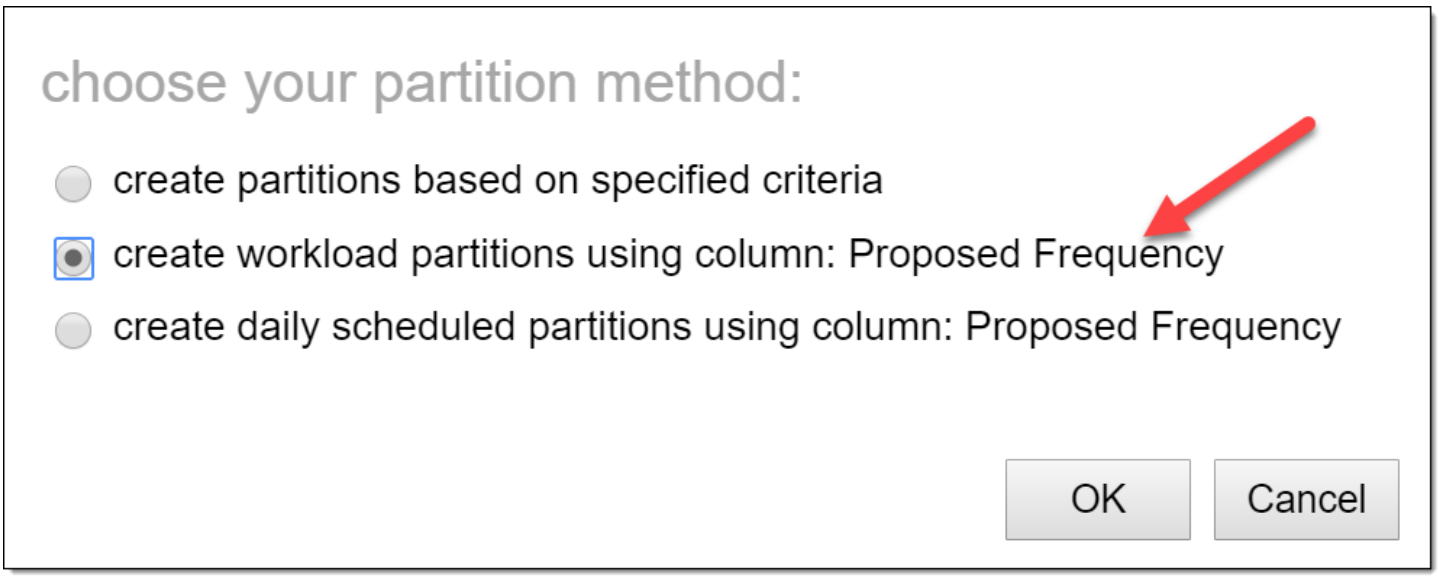
Capacity Planning and Scheduling · https://www.easyterritory.com/documentation/territory-designer/capacity-planning-and-scheduling/workload-partitioning/

Workload partitioning is a powerful tool in EasyTerritory for planners to determine the number of resources needed in a territory or region where locations (accounts, customers, etc.) are repeatedly visited. Workload partitioning takes into account multiple factors:

- How frequently must a site be visited across the schedule period (e.g. every day, week, month, quarter, etc.)?
- How much dwell time is spent on-site?
- What is the driving distance between sites?
- How many hours per day are you targeting for your resources?
- Do you want accounts to always be visited by the same resource or can resources be interchanged?

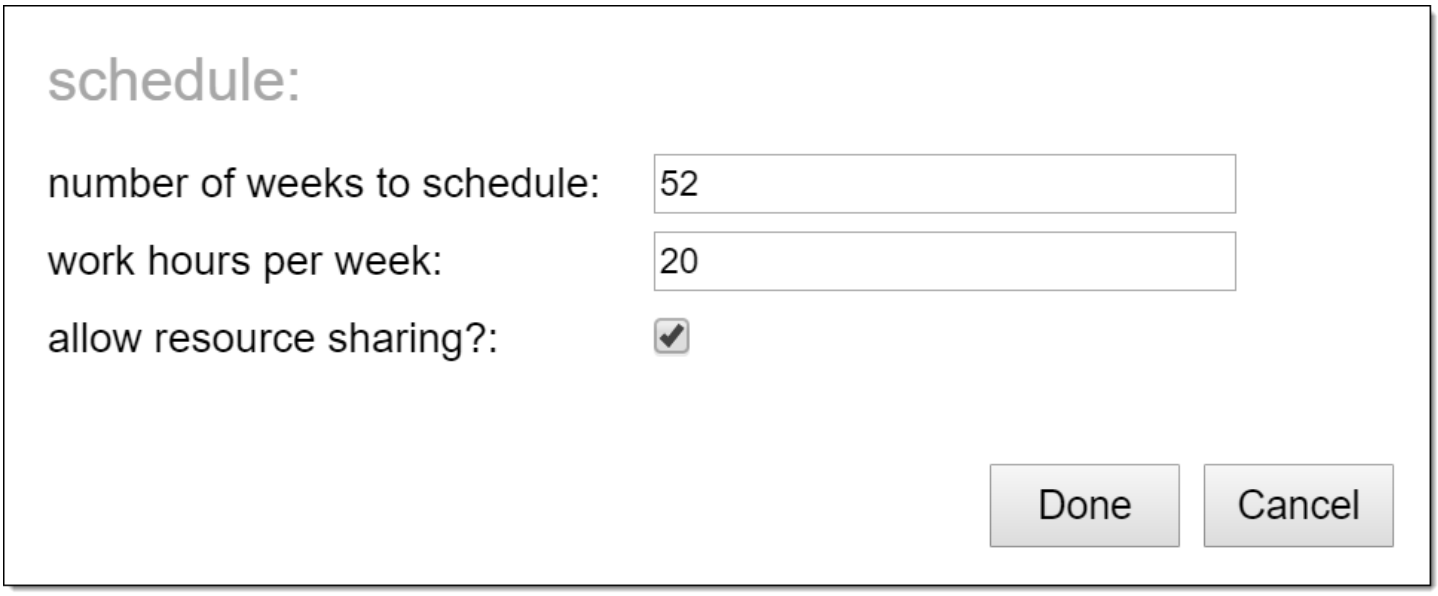
In order to use EasyTerritory's workload partitioning feature, there is only one required data-column which must be included in your dataset and flagged with the appropriate metadata:

- *Visit frequency in weeks* – this is the number of weeks between each visit. To visit a site twice a week, the value is 0.5; for every quarter it is 13. The data column must be flagged with **scheduleweekfreq**
- *On-site time in hours* – this is an optional column and specifies the amount of dwell time per visit for a given site in hours. The data column must be flagged with **routestopdelayhours**. If this column is not provided, then you will be prompted for the average onsite dwell time per stop.



To use the workload partitioning feature, perform the following steps:

- Query the features you want to schedule. This may be done using any of EasyTerritory's powerful querying capabilities (by ctrl-clicking on features; selecting by lasso, envelope or polygon).
- Click on the *partition* link in the results-grid – this will open an options dialog.
- Select the *create workload partitions* option – this opens the scheduling dialog:

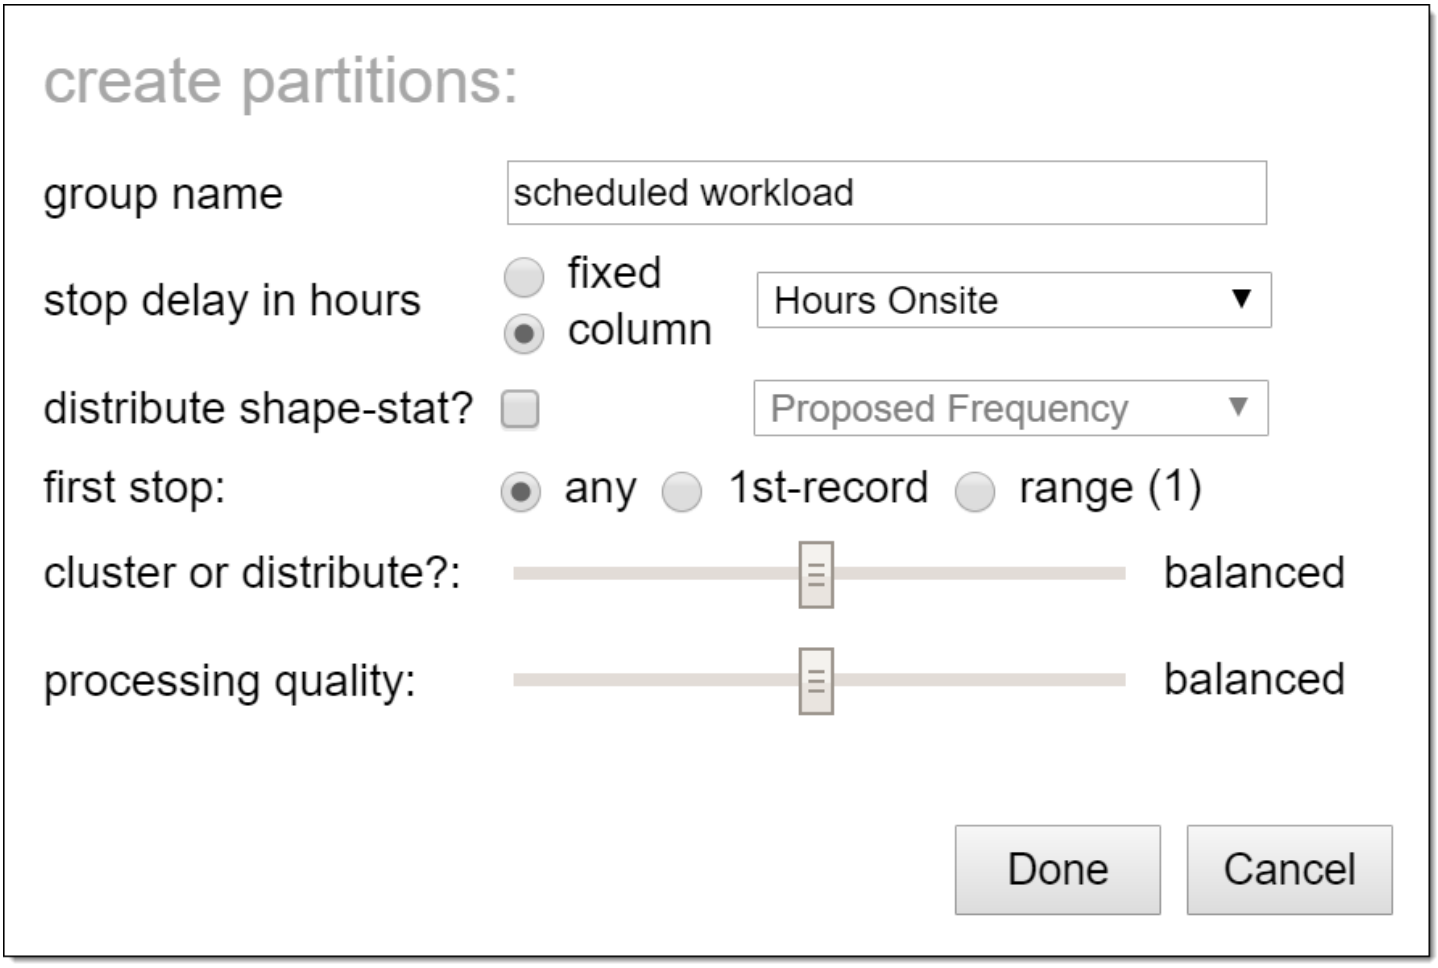


- Enter your scheduling criteria:
  - **Number of weeks to schedule** – this is the number of weeks to analyze – typically this is set to the largest value in your **scheduleweekfreq** column so every stop is included in your

analysis.

- **Work hours per week** – this is the target number of hours you want each resource to work per week.
- **Allow resource sharing** – check this if a location (account, store, etc.) can be visited by any resource. Leave this unchecked if your reps have a relationship with the account and must be a consistent resource on subsequent visits.

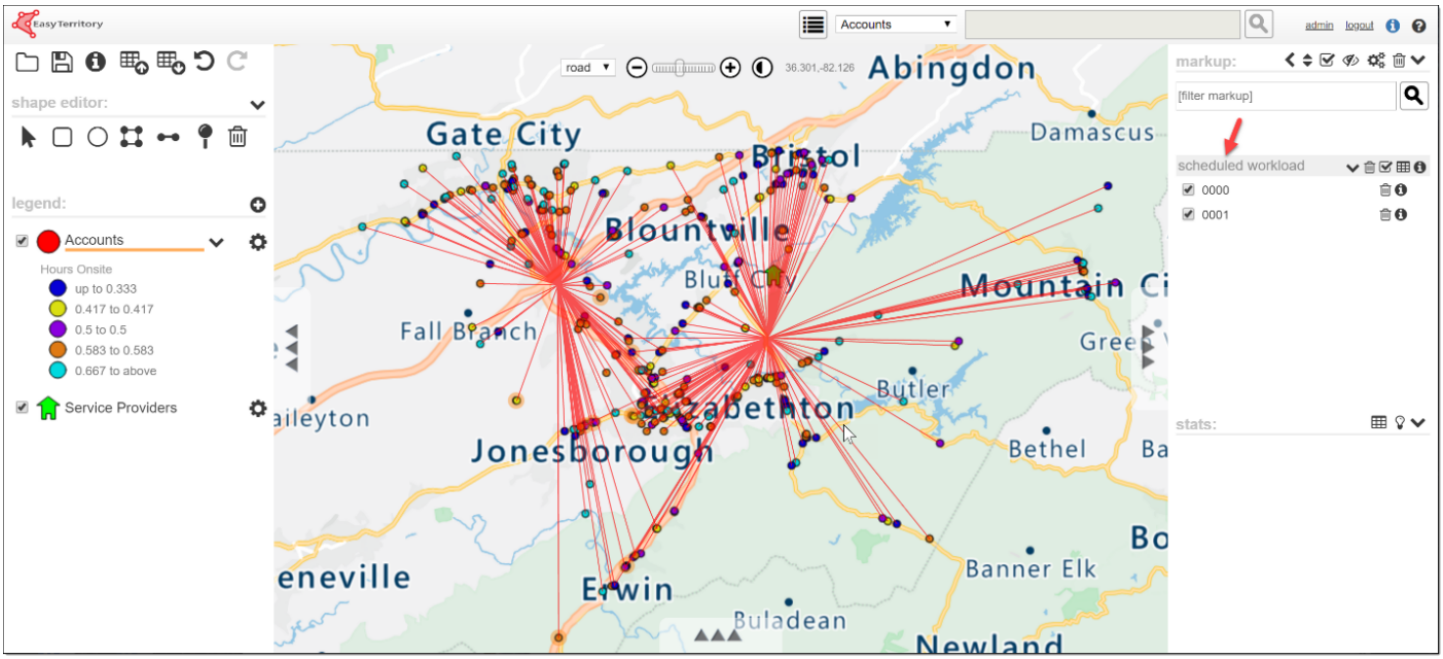
- Press *done* to open the partitioning dialog:



- Enter your partitioning criteria:
  - **Group Name**: this is the name used to group the results in the markup panel
  - **Stop delay in hours**: use your data column flagged with **scheduleweekfreq** or enter an average onsite dwell time in hours used across all stops

- **Distribute shape-stat**: this is an optional column that will be used in the partitioning process to attempt to evenly distribute (e.g. sales volume or store rating)
- **First-stop**: leave this set to *any* by default as a start-stop is not needed in this analysis
- **Cluster or distribute**: there is a conflict between isolated workload clusters with minimal drive times and precise distribution of workload between partitions. Set this slider to your preference. *See summary of clustering vs distribute at the bottom of this article.
- **Processing quality**: this is the amount of time spent trying to optimize the solution.

- Press *done* to create workload partitions:



Based on the number of accounts, frequency of visits, on-site dwell time and expected weekly hours (20), we created two partitions – that is, two resources are needed for this region. From here we might select one of the partitions and use EasyTerritory's scheduling feature to create a daily visit schedule for the resource. Please review our documentation on Scheduling for more information.

*Partitioning: Cluster or Distribute?

One of the powerful features of EasyTerritory is partitioning. Partitioning divides point locations on a map into sets (clusters) of equal workload. A workload is computed by totaling the onsite times and estimated drive-times between each stop in the partition. An ideal partitioning of a set of points is easy to imagine; they are isolated, small and equal in workload. Take for example a set of customer

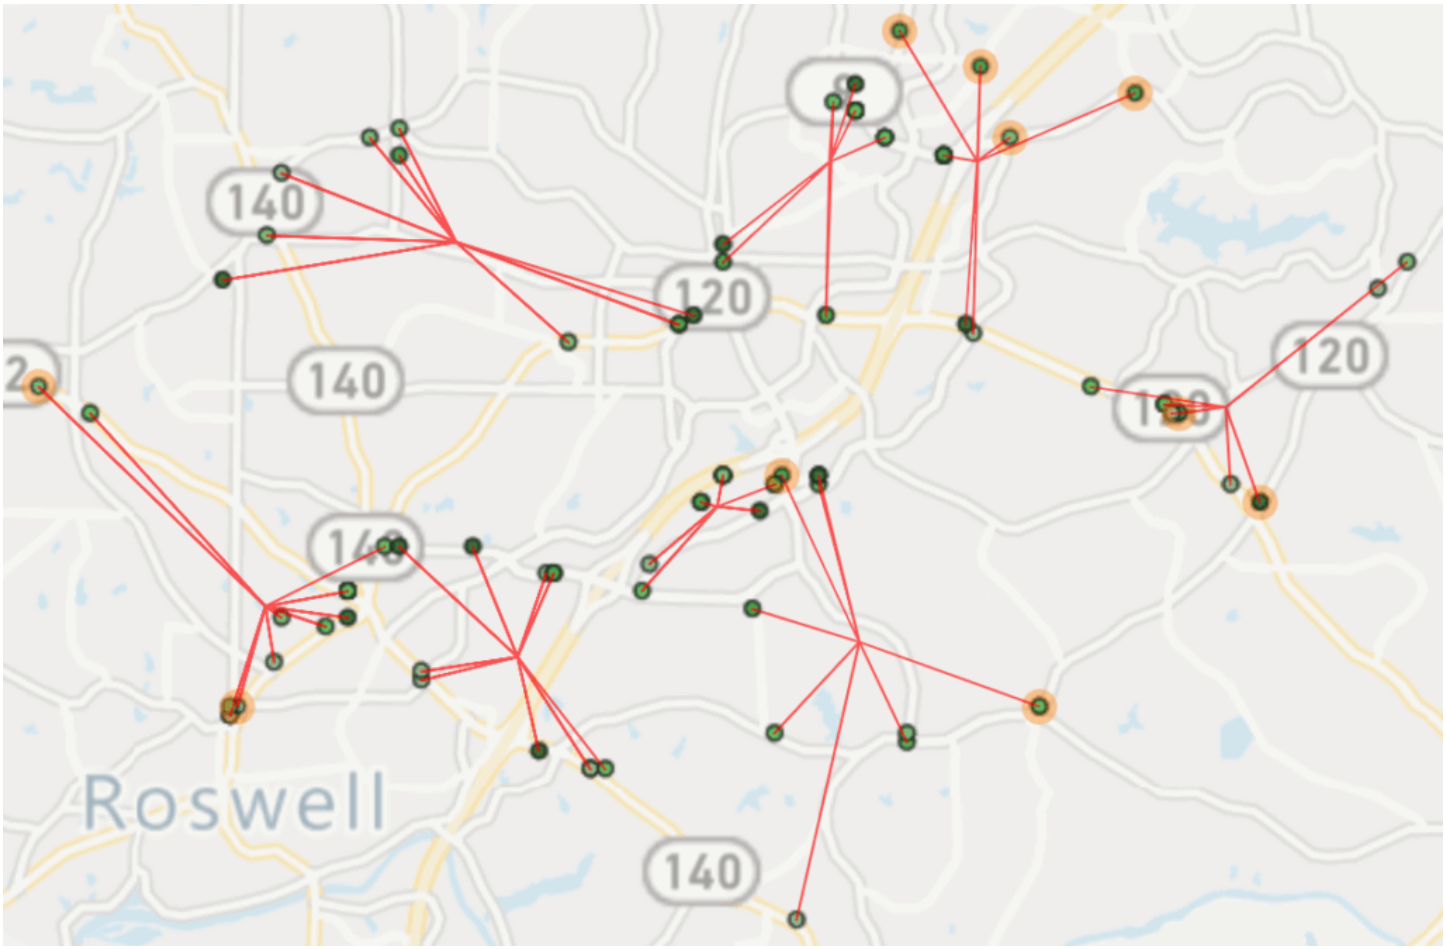
work-orders for servicing point-of-sale equipment. In this example, ideally the service technicians are not crossing each other in transit; they are driving the fewest number of miles and they are spending the same number of hours working and driving. Though easy to imagine, it’s hard to determine without a software like EasyTerritory.

In real-world scenarios, ideal partitioning is rarely achieved. There are almost always compromises. The primary compromise revolves around the tension between small, isolated partitions — we will call these clusters — and equally-distributed workloads. Nearly all scenarios allow for good clustering or equal-distribution, but not both. Some scenarios allow for a good compromise where both are more than acceptable. Other scenarios must be analyzed to determine which compromise is most acceptable. In EasyTerritory, this analysis requires the user to adjust the sliding scale between “cluster” and “distribute” when partitioning.



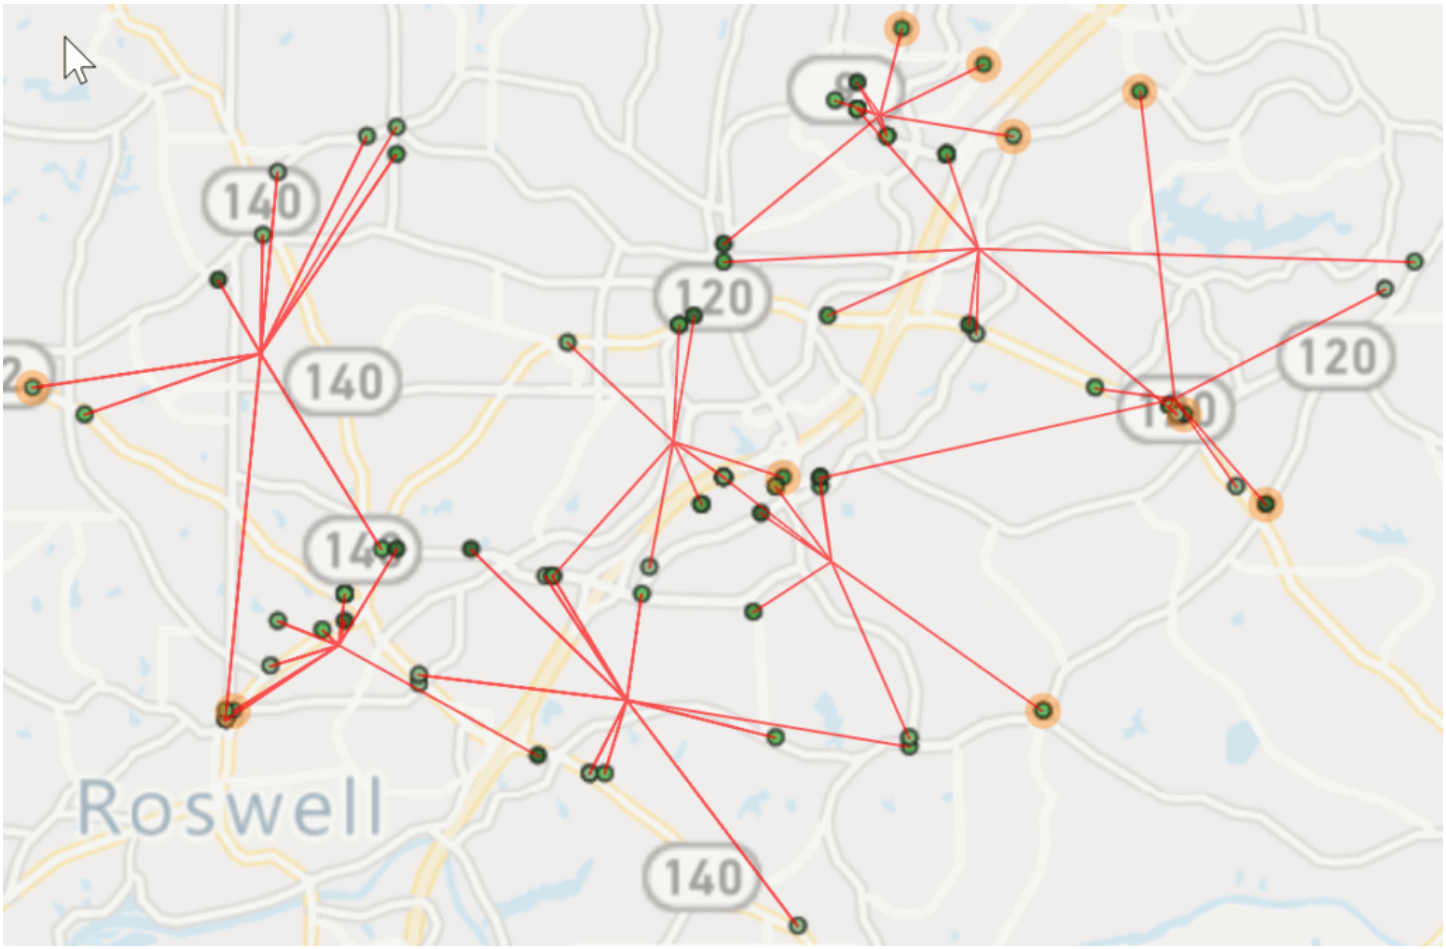
So what happens when you move the slider? That’s a good question and many of our customers leave it balanced in the middle as they are not sure what to set it to. But here’s how it works. If the slider is moved to the left it will bias the partitioning engine towards clustering, resulting in smaller, isolated partitions. Your service techs are not passing each other and their total workload is minimized. Great! But the downside is an increased variance in the cluster’s total workload — total onsite time plus total drive time. One service tech might have an 8-hour day and another a 6-hour day.





If the slider is moved to the right, the partitioning engine is biased towards distribution, resulting in partitions of equal workload. All of your service techs have roughly 7 hours of driving and on-site time. Super! But the downside here is the partitions are larger (more drive-time) and they are more likely to overlap. You look at the results and you wonder why one service tech is assigned to a stop deep inside another technician's route. But that is the mathematical compromise that cannot be escaped in most scenarios.

cluster or distribute?: distribute



So how do planners determine the level of compromise? The challenge is in determining distribution — clustering is simple. For clustering, you simply look at the map in EasyTerritory. If they are less isolated, larger and have more overlap, then you know you are compromising on total drive-times and encroachment and clustering is not optimal. If the partitions are smaller and more isolated, then you know you are maximizing drive-time efficiency and reducing encroachment.

For distribution, the planner is shown CoV in real time as the partitioning engine is running. CoV is the coefficient-of-variation, or sometimes called relative-standard-deviation. The lower the CoV the better. Because it is relative-standard-deviation, (standard-deviation divided by the mean) it always provides a useful value that is more industry-dependent than scenario-dependent. In other words, if a planner finds the partitioning results for scenario A to have a good workload distribution and the CoV is 0.02, then if it is 0.02 in scenario B, the distribution is also likely to be acceptable. The planner will need to try a few scenarios to see how CoV corresponds with what is acceptable in their business practice.

To determine in absolute terms (in hours) each partitioned workload, you have two options. If you bring up the partitions properties in the markup panel in EasyTerritory, you will see a total-time estimate. You can compare these between partitions to determine the quality of distribution. The only way to get a solid comparison is to fine-route each partition using the EasyTerritory routing capabilities. This will give you the most accurate analysis of how well your partitions are distributed.

Finally, the density of stops makes a difference. If your accounts are located miles apart in rural areas (lower density) then you will typically want to bias more towards distribution. Due to the large variance in drive-times, small isolated clusters are almost certainly going to vary greatly in hours of workload. If your scenarios are in metro-areas with many accounts close-by (higher density), then you will likely want to bias more towards clustering. Drive-time is less of a factor and overlap between partitions is more likely.

If you are new to partitioning, a good first step is simply experiment with your data until you get comfortable with the settings and the results. If you find the results are generally good but need a few tweaks, EasyTerritory provides a manual means to address this. Please see our videos on partitioning for more information.