Exeleus Meridian

Easy user guide · Free individual production service · 11 November 2026

1. Your first run

Meridian helps you explore how cyber incidents could affect a network of people, devices and services. You describe the network and its assumptions, then run many simulated scenarios and review their costs.

Start here

Open **crq.exeleus.com** to create a Free account, or choose **Sign In** for an existing account. Complete the eligibility, consent and terms screens using your own correct details.

Free personal accounts are for eligible adults; some school-related accounts need a school-managed route. No invitation or payment details are required. Modeling, basic simulations, results and downloads are included, with **25 accepted runs per month**.

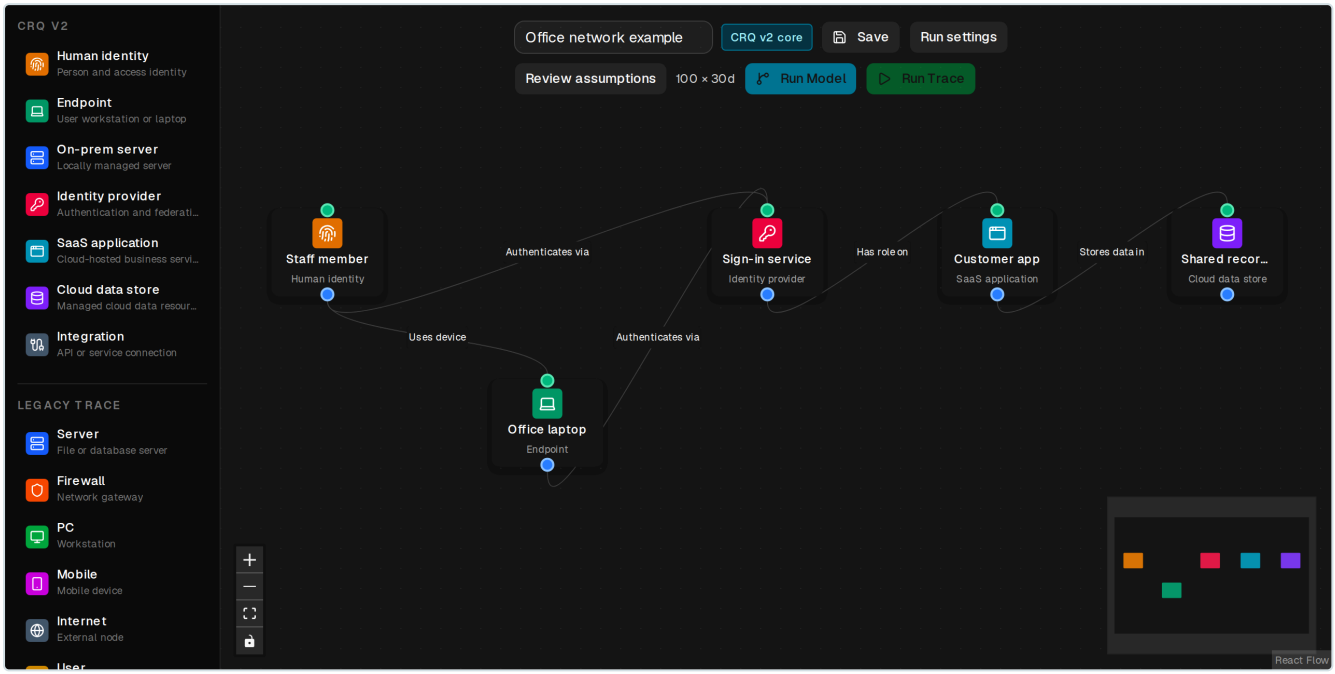
Follow these six steps

1. **Create a model.** Open **Model Topology**, choose a starter model and click **Create model**.
2. **Describe your network.** Check the assets and connections. Give the model a useful name.
3. **Check the costs.** Choose a loss model for each asset. Click **Apply**, then **Save**.
4. **Review the assumptions.** Click **Review assumptions** to inspect the whole network's inputs.
5. **Run the model.** Check **Run settings**, then click **Run Model**.
6. **Read the results.** Wait for **completed**, inspect the costs and download a report.

For your first run, the existing defaults are **100 replications over 30 simulated days**. That is one submitted run containing 100 scenarios—not 100 monthly runs.

The guide's screenshots use an illustrative model and demonstration results. Starter values and example numbers are not calibrated estimates for your organization.

2. Build and save your network



Model Builder: asset palette on the left, network in the center and action buttons at the top.

Start with a template

In **Model Topology**, choose a **Template**, enter a **Model name** and click **Create model**. The model opens in the Builder. A template is a starting point: review its assets, connections and values before running it.

Make the network your own

1. **Add an asset.** Drag an item from the left palette onto the canvas. On a phone, tap **Add agents**, then choose an asset.
2. **Edit an asset.** Select it to open its properties. Set its name and relevant probabilities, values and costs. Click **Apply**.
3. **Connect assets.** Drag from the blue handle at the bottom of one asset to the green handle at the top of another. Select the connection to check its relationship and transmission probability.
4. **Save.** Enter a clear model name and click **Save**. Saving an opened model updates that model.

Use the **CRQ v2** assets for the core workflow: human identity, endpoint, on-prem server, identity provider, SaaS application, cloud data store and integration.

Apply changes the canvas; Save stores the model. Opening a properties panel does not save a change. **Run Model** also saves the current model before submitting it.

Use a larger screen when possible for detailed network editing. Zoom and pan to move around the canvas.

3. Choose costs and review assumptions

Choose how each asset can lose value

Loss option	Use it for	What to enter
Finite asset/data value	A limited data or asset value	Data value and fraction of that value lost per compromised five-minute interval
Recurring operational loss	Ongoing loss while compromised	Operational loss per compromised hour
Finite value + recurring loss	Both limited value and ongoing loss	Both sets of inputs

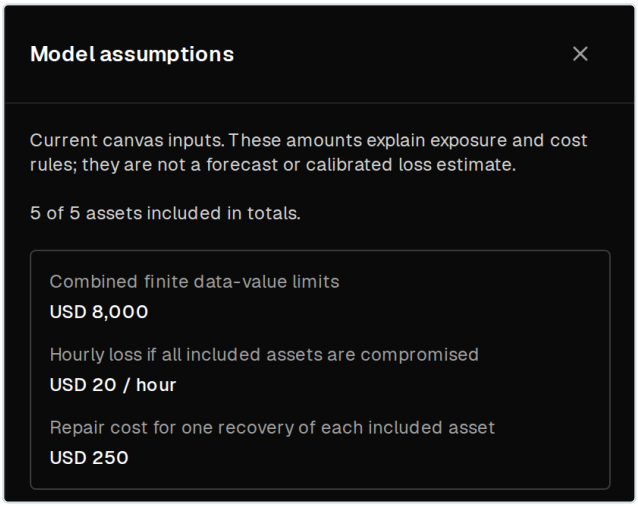
A finite value is the limit within each simulated scenario; recovery does not refill it. Repair cost is charged at every recovery. Keep the same expense out of multiple categories. Input labels use USD; Meridian does not convert currencies.

Why can the total exceed the asset value?

For one asset, a scenario might include:

- **USD 1,000** of finite data loss;
- **USD 100** of operational loss: five compromised hours at USD 20/hour;
- **USD 150** of repairs: three recoveries at USD 50 each.

The total is **USD 1,250**. The USD 1,000 limit caps finite data loss—not the other charges. This is an accounting example, not a forecast.



Review assumptions separates finite limits, hourly exposure and one repair charge per asset.

Review the whole model

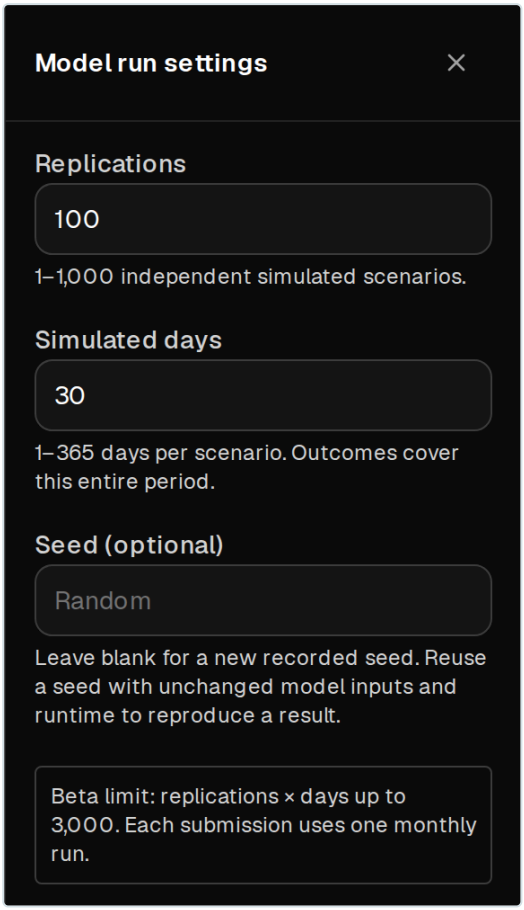
Click **Review assumptions**. Search for an asset and use **Open asset properties** to inspect it. Search does not change the model-wide totals. The amounts shown are input limits and conditional costs, not predicted losses.

Probabilities apply every five minutes. A 2% entry probability per five minutes is much stronger than 2% per day. Use **What these probabilities imply** in an asset's properties to inspect its isolated reference. Older saved assets may still use repeated-value loss until you explicitly change their loss option, Apply and Save.

Fields marked **Budget Screening** belong to the separate screening tool; they do not change core Run Model transitions.

4. Run a simulation

1. Open **Run settings** in the Builder.
2. Set **Replications**: how many simulated scenarios to generate. The default is **100**.
3. Set **Simulated days**: how long each scenario lasts. The default is **30**.
4. Leave **Seed** blank for a new recorded random seed. A seed helps repeat a result when the model, inputs and software version also stay the same.
5. Close Run settings and click **Run Model** once. The model is saved and the job is submitted.
6. Go to **Results** and keep the saved run selected while it finishes.



Run settings: replications, simulated days and optional seed.

Stay within the workload limits

You can choose **1–1,000 replications** and **1–365 days**, provided:

replications × days ≤ 3,000

For example, 100 × 30 is allowed; 100 × 31 is too large. Each accepted submission uses **one** monthly run allowance.

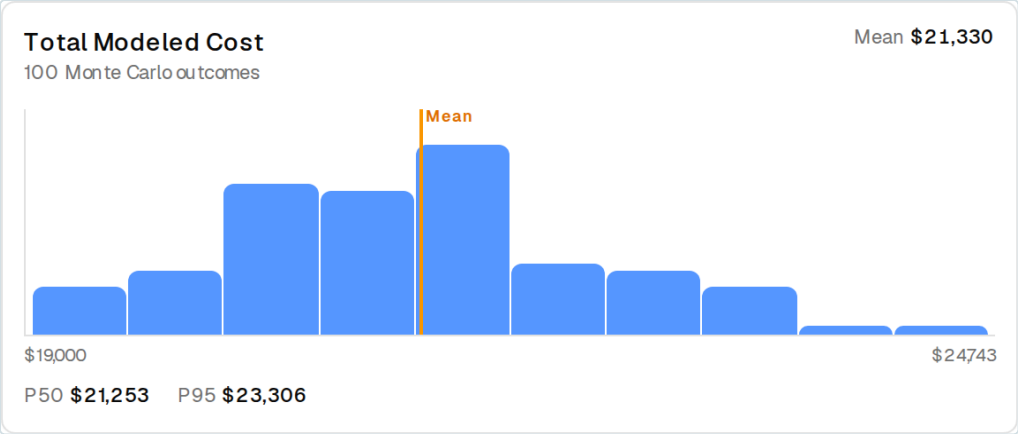
While you wait

A run may wait in the queue while another job finishes. **pending** means it is waiting; **running** means it is executing; **completed** means a saved result is available.

If the page or request appears to time out, check **Results** before submitting again. An accepted run may still be working. You can return to its saved result later.

Use Run Model for this workflow. Run Trace is a separate animation for legacy trace assets.

5. Understand and download results



Example cost chart: Mean, P50 and P95 summarize the simulated scenarios.

Select your run in **Results**. Check its model name, status, number of runs and simulated days before interpreting the amounts. Results cover that recorded period; they are not automatically annualized.

Read the summary

Term	Plain meaning
Mean	Average cost across the simulated scenarios
P50	Median: the middle of the modeled outcomes
P95	A higher-cost point, around the 95th percentile—not a worst-case guarantee
Std Dev	How widely the modeled outcomes vary around their mean
Total Modeled Cost	Data loss, recurring operational loss and repair charges combined

The mean and percentiles describe **individual scenarios**. Do not multiply a mean by 100 simply because you ran 100 replications. Charts show how the scenarios differ; they are not measured incident frequencies for your organization.

Download what you need

- **Export CSV:** open the output in a spreadsheet.
- **Download report:** open a readable HTML report in a browser.
- **Download run archive:** keep the JSON record of submitted inputs and software details with the result.

If a saved result will not load, use **Retry loading results** when shown. If a download fails, retry the download for that same completed run; a new simulation is not needed.

6. Keep your work and recover

Check the model's save status

What the Builder shows	What to do
All model changes saved	The applied model matches the last confirmed save.
Unsaved model changes	Select **Save** to store the current design.
Property edits not applied	Choose **Apply** in the asset or relationship panel, then **Save**.
Saving model…	Wait for confirmation. Edits made during this save still need another save.

The desktop and mobile navigation menus ask before leaving unsaved work. **Cancel** keeps you in the Builder; **OK** leaves it. Save before browser Back/Forward or account-menu actions. Refresh/close warnings depend on your browser.

Keep a model backup

1. Apply property edits, open **Model files** and select **Download model backup**.
2. To reopen a backup, choose its JSON file, review it and select **Open as new draft**. Confirm replacing the current canvas only when ready.
3. Select **Save** to create a new saved model. Existing saved models remain available.

A model backup contains the applied design. Unapplied property edits, results and run settings are separate. Keep the file private if it describes sensitive systems.

After a slow or lost reply

If you edit while **Run Model** is being submitted, those newer edits stay in the Builder. Use **View submitted run** to open the original run without submitting again. Save any newer edits before leaving.

If no submission reply arrives, keep the same Builder and settings open and check **Results** before retrying. Selecting an asset or opening its properties alone does not change that unchanged retry. After refreshing, reopening or changing settings, check existing runs before another submission.

Leaving a page does not cancel work already accepted by the server.

Find an older result

Results search covers the **latest 50 runs**. Keep **Saved link to this run**, or use **Open a run by reference** with the full reference from a saved link or run archive. Searching and reloading read saved results; they do not submit simulations.

7. Help and quick reference

If something looks wrong

Situation	What to do
Run Model is disabled	Apply any pending property edits. Use only CRQ v2 assets and check Run settings for an invalid value or a run over the 3,000 replication-day limit.
Run Trace is disabled	It supports legacy trace assets. Use Run Model for a CRQ v2 network.
The run stays pending	It may be queued. Check the saved run in Results; avoid submitting duplicate copies.
Costs look unexpectedly high	Review five-minute probabilities, the loss option, hourly exposure and repeated repair charges. Check the recorded number of days.
A result differs from an earlier run	Check the model inputs, duration, seed and software version. A blank seed selects a new random seed.
The monthly run allowance is used up	Review existing results until the new month's allowance is available. Contact support if the count looks wrong.
You cannot sign in or access a feature	Use your account email, complete the required account screens and contact support if access is still unavailable.

When you return

Open **Model Topology → Saved models**. Use **Find models** and the sort/page controls, then select **Open model**. Search covers your full saved collection; pages show 20 models. Review changes and **Save**. Completed runs retain their archived inputs. Check **Run settings** again: reopening restores defaults.

Your account and support

Your Free plan works without an invitation. Separately issued free-access grants are optional; they do not raise normal usage limits.

Use **Settings** to download your own account records or withdraw eligibility-processing consent. Withdrawal blocks future feature access while your account-record export remains available; it does not itself erase models or results.

For help or a privacy request, email **admin@exeleus.com**. Include the run reference, approximate time and a short description of the problem. Keep passwords, MFA codes and session tokens private.

A useful habit: review inputs before each run, keep the result's time period in view, and save its report and archive together. Treat the outputs as results of your configured scenario—not a guaranteed future loss.

Open Meridian: **crq.exeleus.com**