Interact DC User Guide

Description	How to use Interact DC to model and improve data centre efficiency.
Copyright	© Interact DC

Contents

Interact DC User Guide 6

- The shape of the work 6
- Who this guide is for 6
- Getting help 6

Your account 7

- Getting an account 7
- Signing in 7
- Setting up your profile 8
- Email notifications 8
- Organisations 8
- Who can see your work 9
- Inviting colleagues 9
- Leaving an organisation 9

Run your first simulation 11

- 1. Create the simulation 11
- 2. Bring in your estate 11
- 3. Check what we matched 14
- 4. Choose what to compare against 14
- 5. Set your own figures 15
- 6. Run it 15
- 7. Read the result 16
- 8. Generate a report 17

Specifications 19

- The three pieces of a simulation 19
- Building one 19
- Hardware configurations 19
- How matching works 20
- Correcting a match 20
- Editing lines 22
- What if the hardware is not in the catalogue? 23

Importing from CSV 24

1. Upload 24
2. Map your columns 24
3. Preview 24
4. Confirm 24
File requirements 25
Common problems 25

Targets and runs 26

Adding targets 26
Simulation types 27
Statuses 27
Cost parameters 28
Submitting a run 28
Run history 29
Duplicating for what-if work 29
Finding a simulation again 29
Folders 30

Reading your results 31

The five measures 31
Servers left out of a run 31
What the colours mean 31
Power and energy 32
Efficiency grades 32
How far to trust it 33
Comparing targets 33
When part of a run fails 34
Exporting 34

EU compliance reports 35

Building a report 35
Reading the results 36
The label 36
The PDF 37

Reports 38

Generating one 38
Choosing what goes in 38

Branding 38
Before you send it 39
Sharing 39

Credits and billing **40**

What a run costs 40
Seeing your balance 40
Personal allocations 40
Refunds 40
Buying more 41
Receipts and credit notes 41
Running out mid-project 42

Administering your organisation **43**

Dashboard 43
Members 43
Credits 44
Settings 45
Branding 45
Server SKUs 45
Activity log 46
Analytics 46

Server SKUs **47**

The SKU list 47
Defining one 48
Importing a spreadsheet 48
Using a SKU 49
Editing, retiring and deleting 49
Where your reference appears 50

Locations and estate **51**

Locations 51
Server groups 51
Bulk upload 52
Hardware configurations 52

Common tasks **53**

Glossary **56**

Getting help **58**

Inside the app 58

Troubleshooting 58

Reporting a wrong result 59

Contact us 59

Terms, privacy and acceptable use 60

Something in this guide is wrong 60

Interact DC User Guide

Interact DC models the servers you run today against the hardware you are considering, and tells you what changes: energy, carbon, floor space, and cost.

Nothing you do here touches your production systems. You describe your estate, we match it against a catalogue of real hardware configurations, and the analysis runs on our side.

The shape of the work

Every piece of analysis follows the same four steps.

1. **Describe what you run.** A *specification* is your current estate — one line per server configuration. Upload a CSV or enter servers by hand.
2. **Choose what to compare it against.** A *target* is hardware you are considering. Add as many as you like and see them side by side.
3. **Run it.** Each run costs credits, and you are shown the cost before anything is spent.
4. **Read the result.** Energy, carbon, consolidation and cost, with a report you can export as a PDF.

> **There is a guided tour inside the app**
>
> Open the profile menu in the top right and choose **Run your first simulation**. It walks you through the same path on the real screens.

Who this guide is for

It assumes no prior experience with the platform. Whether you manage infrastructure and run your own analyses, prepare proposals for clients, or quantify carbon reduction for a sustainability report, this covers what you need to get started.

Getting help

If something here does not match what you see, or you are stuck, see Getting help.

GETTING STARTED

Your account

Getting an account

Almost everybody arrives by invitation. Your organisation's administrator, or we on their behalf, create the account and you are sent a welcome email.

That email carries a **Choose your password** link. Follow it, set a password, and you are signed in and inside your organisation. No password is ever emailed to you, and we will never ask you for one.

The link works once and expires. If it has expired by the time you get to it, use **Forgot your password?** on the sign-in page to be sent a fresh one — your account already exists, so nothing is lost.

Creating one yourself

If sign-up is open, you can create an account yourself. Follow the link from interactdc.com, or choose **Create account** at the top of the sign-in page. You are asked for three things — your email address, a password, and your agreement to the terms — and you are signed in straight away. There is no waiting on a confirmation email to get started.

If the address already has an account, the page says so. Choose **Sign in** instead, or **Forgot your password?** if you don't remember the password.

You land as the administrator of your own personal organisation, with no credits. You can build a simulation, upload your estate and map its columns without paying anything; credits are only spent when you *run* a simulation. See Credits.

We do send a confirmation email, and you need to follow it before you can buy credits — your VAT receipt goes to that address, so it has to be one that reaches you. Nothing else is held up by it.

If sign-up happens to be closed, the page tells you so and offers a way to get in touch; an invitation is then the way in — ask your administrator, or get in touch.

If you already have an account and your colleagues are using an organisation you should be part of, you can ask to join it. An administrator of that organisation approves or declines the request, and you are emailed the outcome.

Signing in

Interact DC uses your email address to sign in — there is no separate username. Always use the address your account was created with.

If you have forgotten your password, use **Forgot your password?** on the sign-in page. You will be emailed a link to set a new one.

> **The link in your welcome email is a password reset**
>
> Setting your password through it signs you in, and it is the same mechanism as **Forgot your password?**. So if you never used the original link, or it expired, asking for a reset gets you to exactly the same place.

Setting up your profile

Open the profile menu in the top right and go to account settings. Add your first and last name so colleagues can tell who did what — otherwise you appear to them only as an email address. Sign-up does not ask for your name, so this is where it gets set.

The same page is where you manage your email addresses and change your password.

Email notifications

Account settings also carries an **Email Notifications** panel, where you choose which mail we send you:

Switch	What it controls
Simulation run results	One email when a run you submitted finishes, whether it succeeded or failed. One per run, not one per target
Compliance report results	One email when an EU compliance report you created finishes, whether it succeeded or failed. One per report, not one per server
Low credit warnings	Sent when your organisation's balance falls below its alert threshold. Administrators only
Requests to join	Sent when somebody asks to join your organisation. Administrators of a shared organisation only

You will only be offered the switches that apply to you, so the panel is shorter for a member than for an administrator.

> **Some email cannot be switched off**
>
> Security mail — password resets and address confirmations — account setup, and billing records always arrive. A VAT receipt is a legal document, and a welcome email carries the only link that makes the account usable, so neither is optional.

Organisations

Your account belongs to an **organisation**. The organisation is what owns your work: simulations, uploaded estates, reports, and the credit balance everything is charged against. You are always working inside it, and you only ever see data that belongs to it.

Three roles exist within an organisation:

Role	Can do
Viewer	Read shared work and produce reports from it. No creating, running or spending
Member	Everything a viewer can, plus create and run simulations and spend credits
Admin	Everything a member can, plus invite people, manage roles, and manage credits

If you signed up on your own, you have a personal organisation of your own and you are its admin.

Who can see your work

Every simulation has a visibility setting, chosen when you create it and changeable afterwards by whoever created it.

- **Private** — only you can see it. This is the default.
- **Shared** — everyone in your organisation can see it.

Choose **Shared** if a colleague needs to read the report. Shared simulations can be read, duplicated and deleted by other people in your organisation, but only the person who created one can edit its specification, targets and settings.

Organisations are sealed from one another. Nobody outside your organisation can see your estate, your simulations or your reports.

Inviting colleagues

Organisation admins can invite people from the members page. An invited person gets an email with a link to set their password and join.

Admins can also change someone's role, reset their password, and suspend an account. A suspended account cannot sign in until it is reactivated. The whole of that is covered in Administering your organisation.

Leaving an organisation

You can leave an organisation you no longer work with, from the organisation page. Your account continues to exist — you land in a personal organisation of your own, as though you had signed up by yourself.

What you leave behind stays behind. Simulations, uploaded estates and reports belong to the organisation rather than to you, so they remain there for your former colleagues, and you will no longer be able to see them. Take copies of any reports you need first.

The owner of an organisation cannot leave it while they own it. Make another administrator the owner first — see Administering your organisation.

If you want the account itself closed rather than moved, ask your administrator, or get in touch.

Run your first simulation

This is the whole path, from an empty list to a finished report. It takes about ten minutes if you have a CSV of your servers to hand.

By the end you will have used every major feature of the Simulations module.

1. Create the simulation

From **Simulations**, choose **New simulation**.

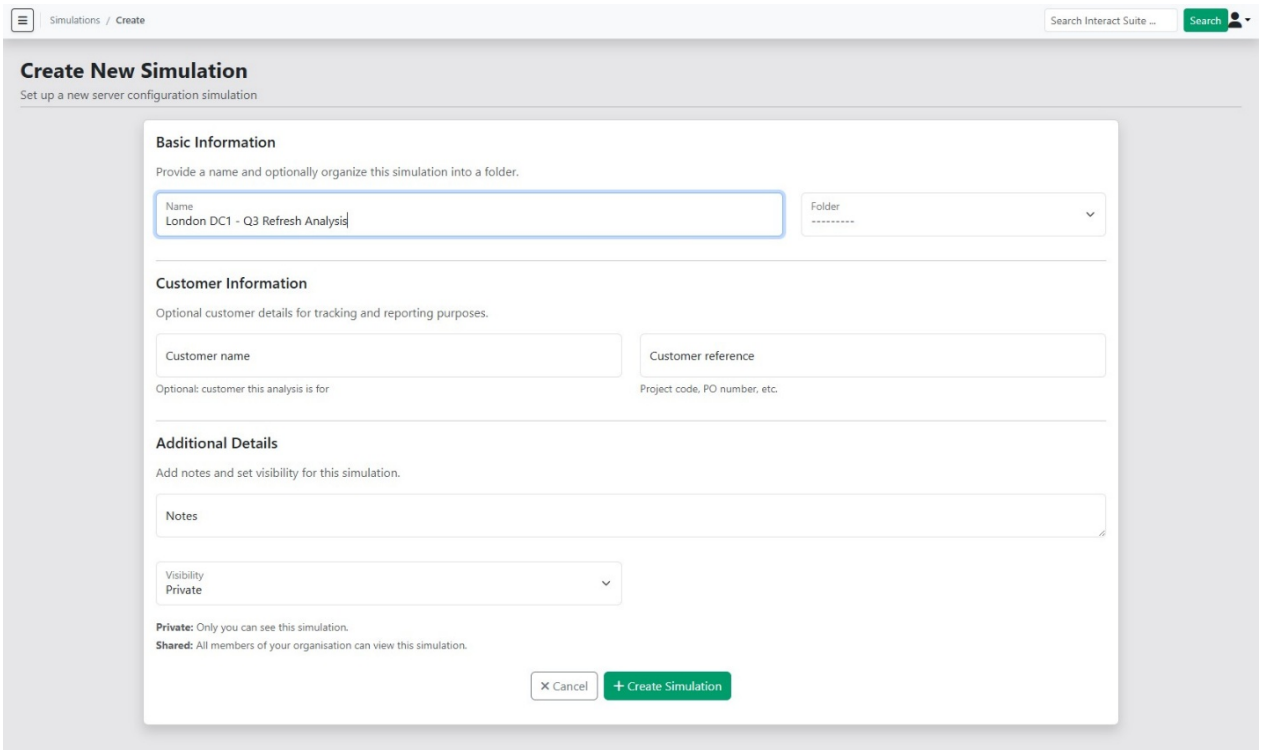
Name it. Something you will recognise in six months: the customer, the site, or the question you are answering — "London DC1 — Q3 refresh analysis" rather than "test 2".

Add customer details. Customer name, reference and notes are all optional, but they carry through onto the finished report and make simulations easier to search for later. Fill them in if this is going to a customer.

Choose a folder if you want to keep related work together.

Set visibility. *Private* keeps it to you; *Shared* lets your colleagues read it. See Your account.

Choose **Create Simulation**. You land on the simulation detail page, which is your command centre for everything that follows.



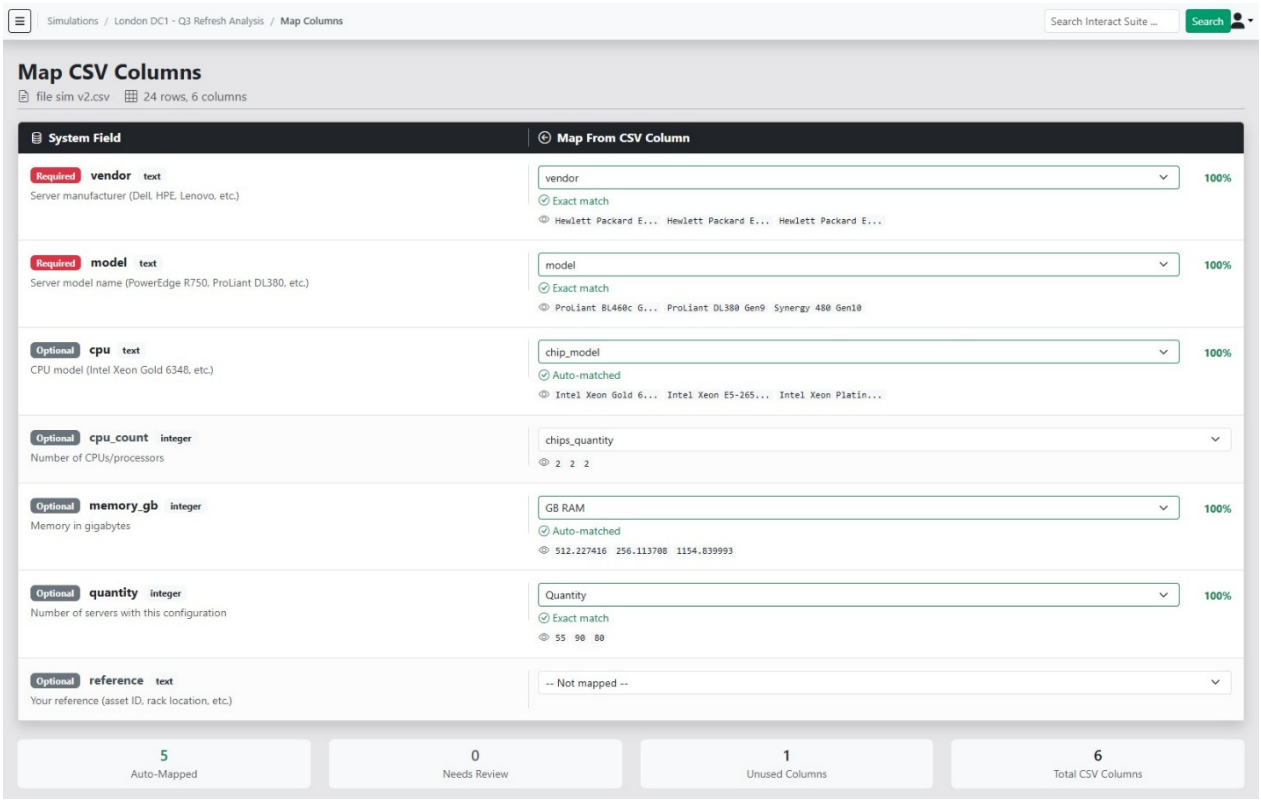
2. Bring in your estate

You need to tell us what servers you run today. There are three ways.

Upload a CSV

The usual choice for anything above a handful of servers. In the **Source Specification** section, choose **Upload CSV** and select your file.

Nothing is imported when you upload. You are shown which columns we found and asked to confirm them first, so a wrong guess costs you nothing.



Map each of your headers to the field it holds. Vendor and model are required; CPU, CPU count, memory, quantity and reference are all optional but improve the match. We suggest mappings for headers we recognise.

Choose **Confirm Mapping and Import**. Matching runs in the background and you are shown its progress.

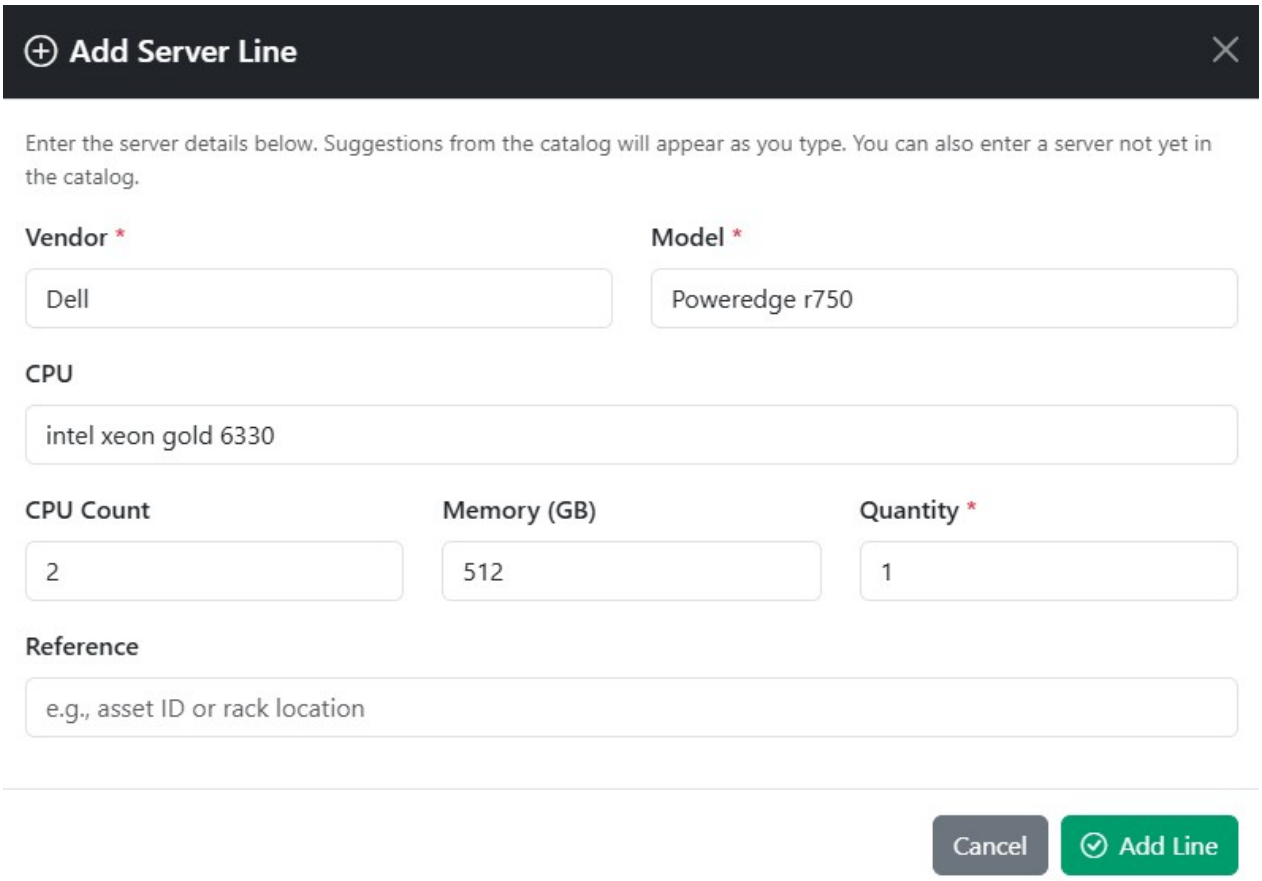
> **Use the template if you are starting from scratch**
>
> The CSV template linked on the simulation page is already in the right shape, with the expected headers in place.

The full detail is in Importing from CSV.

Add lines by hand

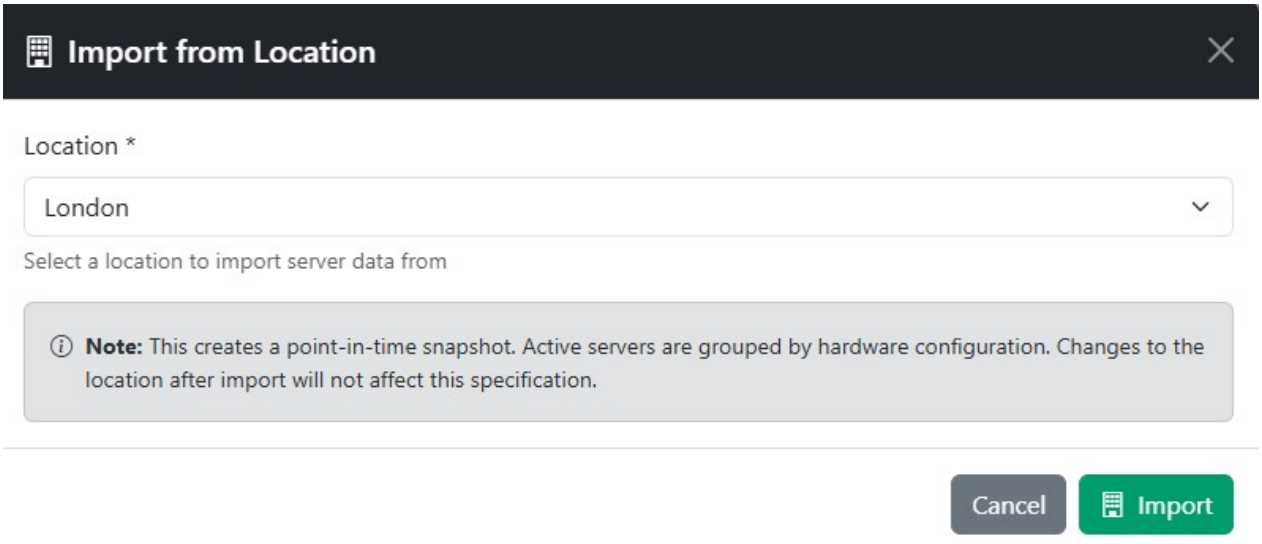
For small estates or one-off additions, choose **Add Data** and then **Add Server Line**. Enter vendor, model, CPU, CPU count, memory, quantity and an optional reference.

Suggestions appear from the catalogue as you type, and an indicator tells you whether we recognise the configuration before you save it.



Import from a location

If your organisation uses the Estate module and already has servers registered under a location, choose **Import from Location** and pick it. Every server there is copied into your specification, so you do not enter anything twice.

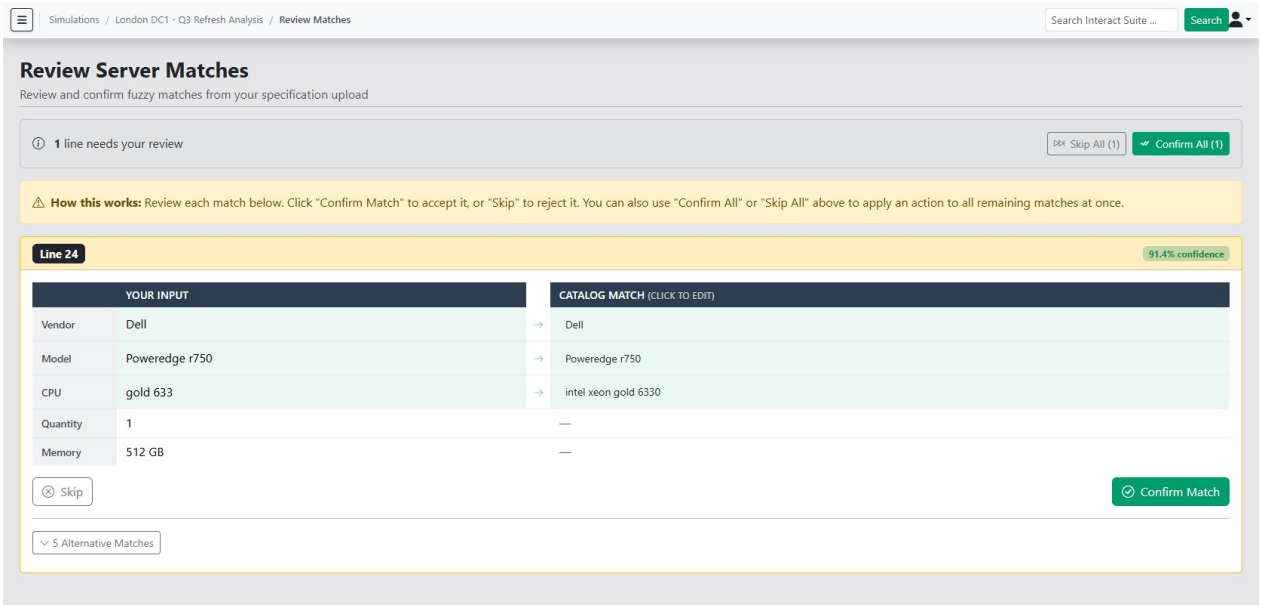


See Locations and estate for what a location is.

3. Check what we matched

Each line is matched against our hardware catalogue. Matching is fuzzy, because inventories rarely spell hardware the way vendors do.

Lines that matched exactly are marked validated and need nothing from you. Lines that matched approximately are queued for review, and a banner appears on the simulation page. Choose **Review Matches** to work through them.



Each card shows your original values, our suggested match, and how confident we are:

- **Confirm** if the suggestion is right.
- **Select Alternative** to browse other catalogue entries if it is wrong.
- **Skip** if the server is not in the catalogue. Skipped lines are left out of the analysis entirely.

Once every line is dealt with, the simulation moves to **Ready**.

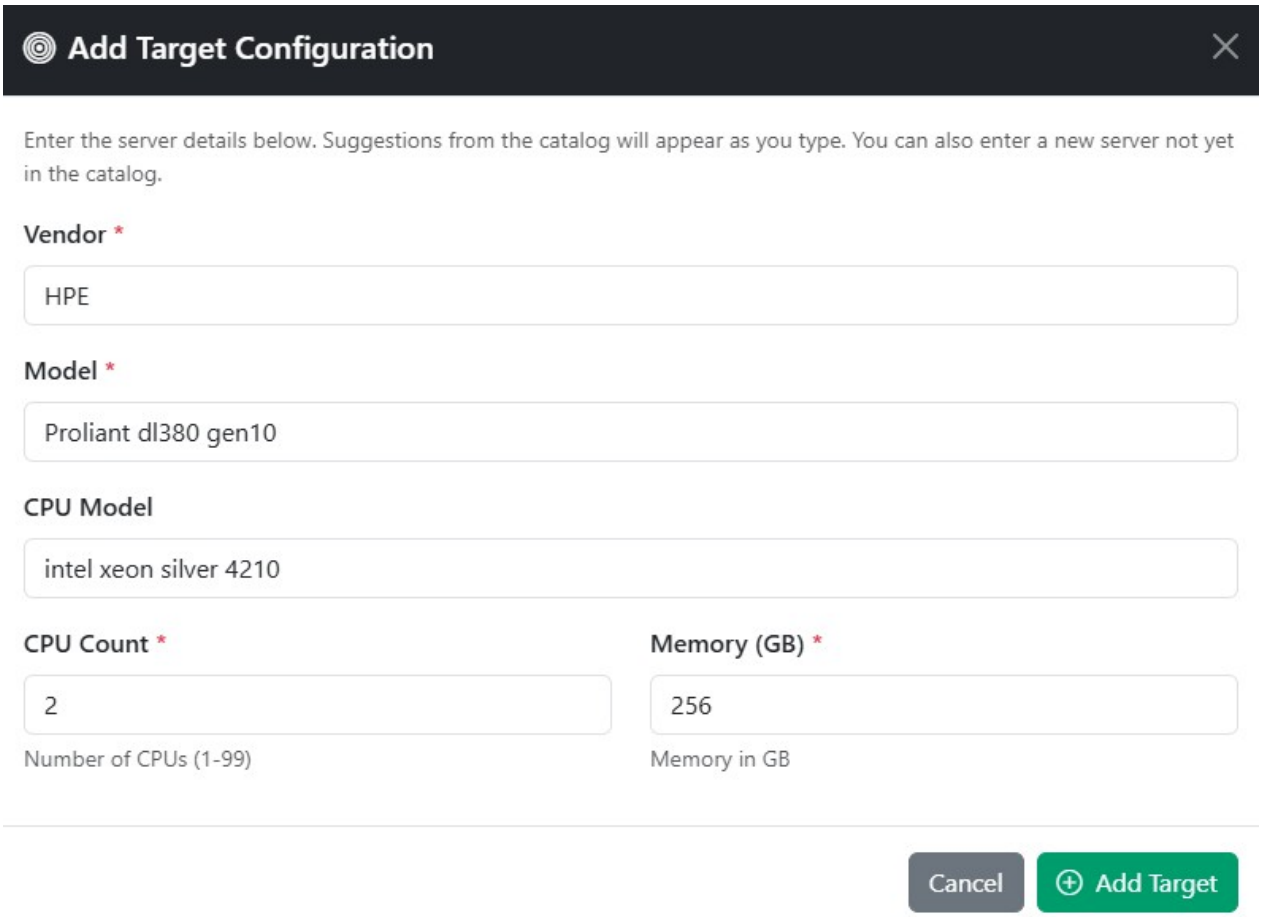
> **Confirm All when the suggestions look sound**
>
> With a long list of reasonable suggestions, **Confirm All** accepts the lot. You can still come back and correct individual lines afterwards.

See Specifications for how to correct a match later.

4. Choose what to compare against

Scroll to **Target Configurations** and choose **Add Target**. Enter the vendor, model and CPU model — suggestions come from the catalogue as you type — then set CPU count and memory.

Once the target resolves, its efficiency grade is shown, so you can see how it compares to your current hardware before you spend anything.



Add as many targets as you like. Each is analysed independently and they appear side by side in the results.

> **Each target costs one credit per run**
>
> Five targets means a run costs five credits. See Credits and billing.

5. Set your own figures

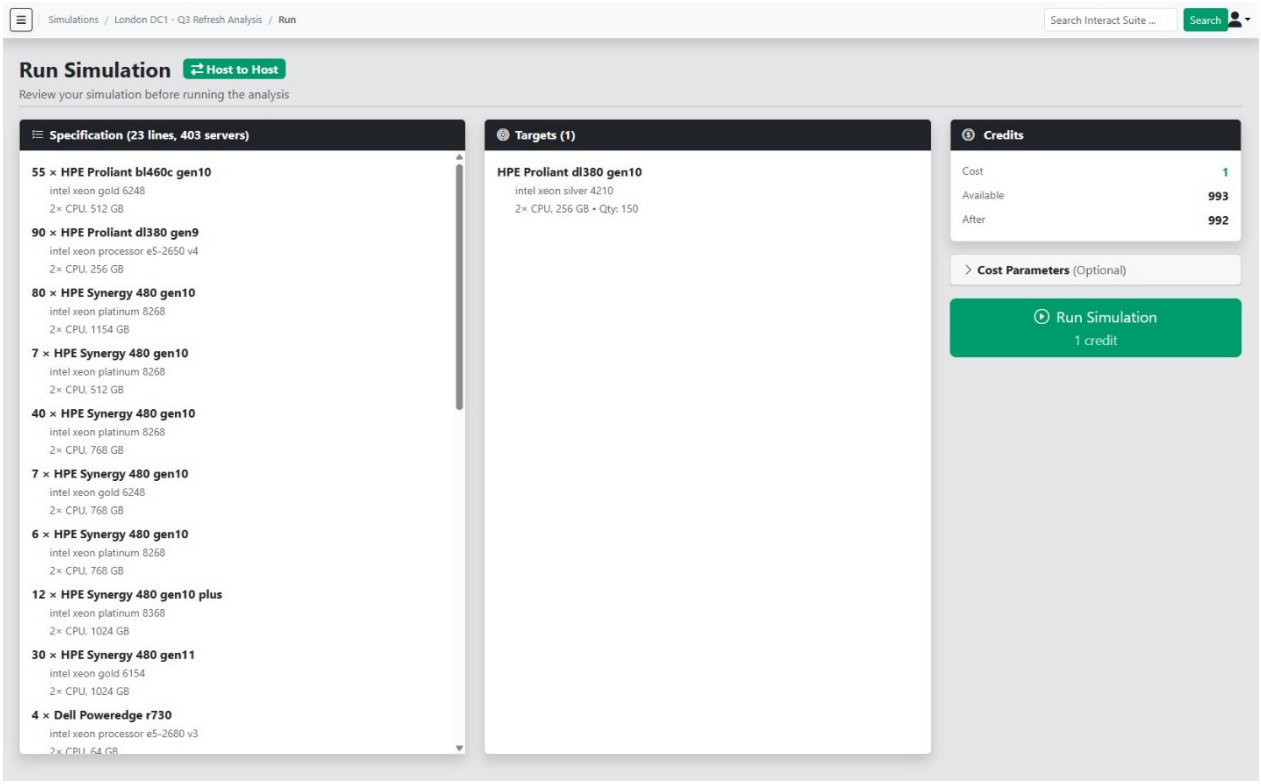
Energy price, maintenance cost per server, and licensing cost per core. Anything you leave blank uses our platform assumption, and the report states plainly which figures were assumed rather than supplied.

Supplying your own energy price is the single change that most improves the accuracy of the cost result.

6. Run it

With a validated specification and at least one target, the simulation is **Ready**.

Choose the simulation type — see Targets and runs — then **Run Simulation**. You are taken to a confirmation screen showing how many source servers and targets will be analysed, what it will cost, and what your balance is.

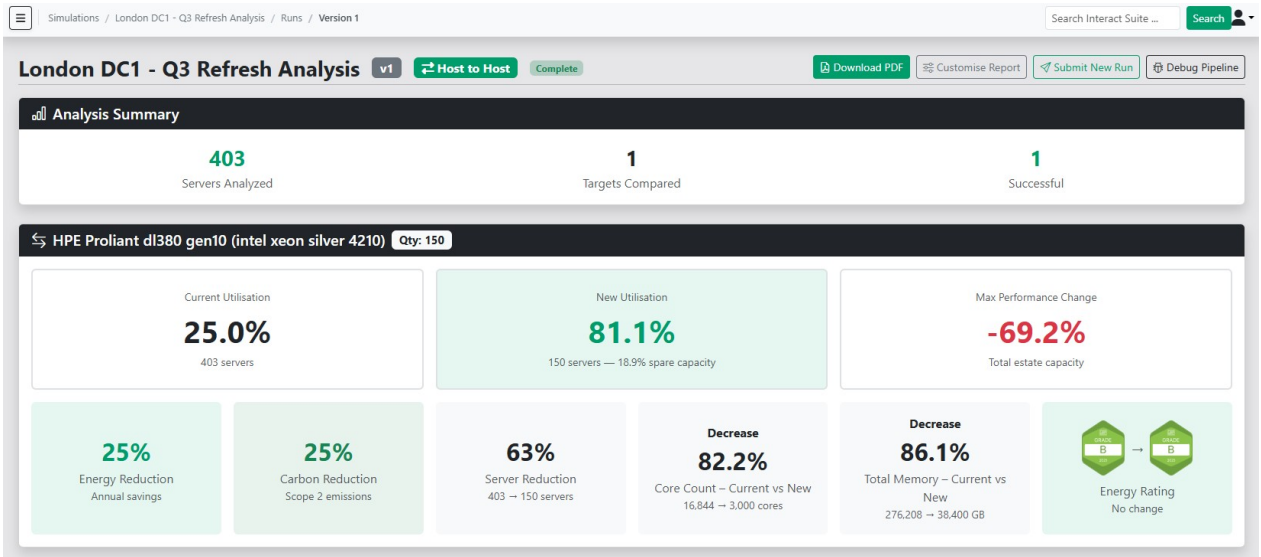


Nothing is charged until you confirm. Credits come off at that point and the analysis starts.

A run usually takes a few minutes. You do not need to stay on the page — the simulation updates itself when the result arrives.

7. Read the result

Results appear on the simulation page and in the run history. For each target you get annual energy reduction, carbon reduction, server consolidation, the new efficiency grade, and operational cost reduction if you set cost parameters — with absolute figures alongside the percentages where we have them.

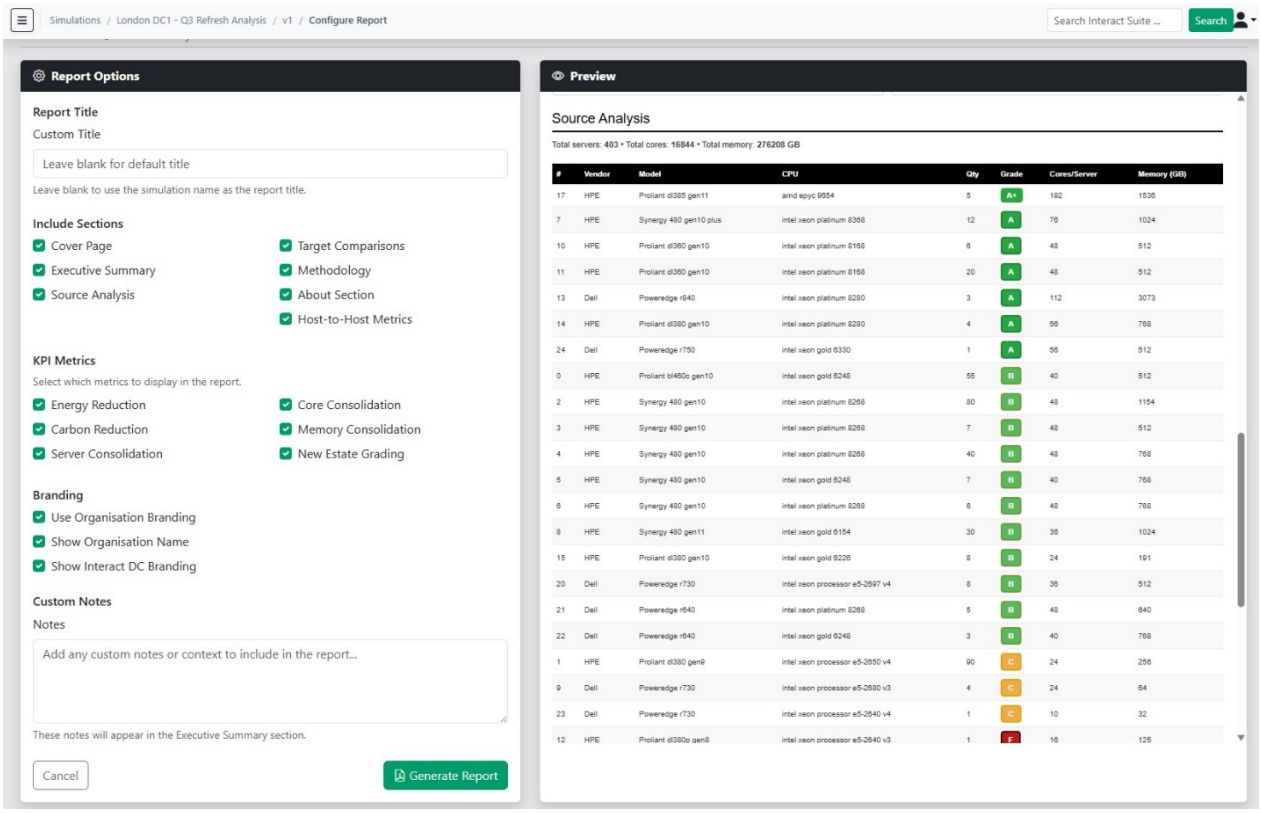


> **A partly failed run still completes**
>
> If some targets fail and others succeed, the run is marked complete and the failed targets are refunded automatically. The error is shown on the target itself.

The full detail is in Reading your results.

8. Generate a report

From the run, choose **Generate Report**. You can pick which sections to include — cover page, executive summary, source analysis, target comparisons, methodology, about — toggle individual metrics on or off, and add your own notes.



If your organisation has set up branding, the report uses it automatically. Choose **Generate**, then download the PDF and send it on.

> **Tip**
>
> **Quick Generate** produces a report with the default settings and no configuration step.

See Reports for what goes into each section.

The same walkthrough runs inside the app

Open the profile menu in the top right and choose **Run your first simulation**. It points at the real screens as you go.

BUILDING A SIMULATION

Specifications

A **specification** is your current estate as Interact DC understands it: one line per server configuration, with the counts, CPUs, memory and any other detail you supplied.

Each line is a **specification line**, and represents one or more identical servers — fifty identical R750s are one line with a quantity of fifty, not fifty lines.

The three pieces of a simulation

It helps to keep these apart, because the app uses the words precisely.

Specification — what you run today. Built by CSV upload, by hand, or by importing a location.

Targets — the hardware you are considering instead. See Targets and runs.

Runs — one execution of the analysis, comparing the specification against every target. A simulation keeps all of its runs.

Building one

There are three ways in:

- **CSV upload** — the usual route for anything above a handful of servers. See Importing from CSV.
- **Manual entry** — add individual servers. Sensible for small deployments or for filling a gap after an import.
- **Location import** — bring in an entire site at once, if your organisation uses the Estate module. See Locations and estate.

Hardware configurations

Every line is matched to a **hardware configuration** in our catalogue: one server specification, identified by vendor, model, CPU model, CPU count and memory.

Configurations carry a pre-calculated efficiency grade, and that grade is what every energy figure in your results is built from. Servers sharing a specification point at a single catalogue entry rather than each holding their own copy — which is why correcting a match once fixes it everywhere it appears.

A configuration we have not graded before is sent for grading as soon as you confirm its match. Grading happens in the background, and the line shows a clock in its status column until it lands — hover it for the detail. There is nothing to wait for: a run grades anything still outstanding before it starts, so a pending grade never reaches your results.

Lines matched only approximately are **not** graded until you confirm them on the review screen, and neither are lines whose match you rejected. Grading a guess would tell you about a server you are about to correct, so we wait for your decision. Nothing is held up by this: a simulation cannot run until its matches are reviewed anyway.

How matching works

Matching runs in stages. We check for a known alias first, then an exact match on vendor, model and CPU, then a fuzzy match if neither lands.

The fuzziness is deliberate: inventories rarely record hardware the way vendors name it, and an exact-match-only approach would reject most real data.

Each match carries a **confidence score**. Lines we are confident about are matched silently and marked validated. Lines we are unsure about are flagged for you to confirm.

A line that differs from our catalogue entry only in how it was written is not flagged: `Dell PowerEdge R7615 / AMD EPYC 9534` is `PowerEdge R7615 / AMD Epyc 9534`, whatever the score. The vendor repeated in the model, capitals, spacing and a CPU's clock speed are all ignored. A different suffix is never ignored — an R740 is not an R740xd.

A CPU is always matched to the model number you wrote: a Gold 6238 is only ever matched to a Gold 6238, never to a close-sounding neighbour. Notes some inventory tools add, such as `(2.10 GHz / 22-core)` or `(8 Cores)`, are ignored. If we do not hold that model number, the CPU is left for you to pick rather than guessed.

HPE servers can be written the short way: `BL460c G9`, `DL380 G10` and `Synergy 480 G10 +` are read as `ProLiant BL460c Gen9`, `ProLiant DL380 Gen10` and `Synergy 480 Gen10 Plus`.

Correcting a match

One card per distinct entry, not per line. If your file names the same server five hundred times, we matched it once and the review screen asks you about it once — the card says how many lines it stands for, and confirming or skipping settles all of them. A 1,000-line file usually comes down to a few dozen decisions.

Lines are grouped on what you wrote: vendor, model and CPU, ignoring case and extra spaces or tabs. The same server with two different processors is two cards, because they are two different matches.

A decision carries to the other cards. When you confirm a card, the model you confirmed is filled in on every other card in the import with the same vendor and model, and the CPU on every card with the same CPU. A card that then has nothing left to decide settles on its own. A card that still has an open question stays, with the decided field marked *As decided on line N*, so you can see where it came from and change it if that decision was wrong.

A group covering more than 25 lines says so above the buttons, with the entry it is about to settle. Worth reading before you press it — one click really does settle five hundred rows.

Whatever you decide applies to **this import only**. Nothing is remembered for the next file you upload.

Open the flagged line. The right-hand column is what we matched it to, and every box in it can be changed.

Click a box — vendor, model, CPU or GPU — and it opens with what we considered for that field, ranked, with a confidence figure against each. Ours is at the top marked **our match**, so you can always get back to it. Pick one and it fills the box; nothing is decided yet. You can also just type: each box searches the catalogue as you go, and you are free to enter something it does not contain.

Being able to change the boxes separately matters more than it looks. The model and the CPU are matched independently, and CPU names differ by a single letter — a `6338` and a `6338N` are different parts with different efficiency. Setting the right model and the right CPU is one decision, made in two boxes.

Confirm Match is what commits, applying whatever the boxes say. Until you press it, nothing about the line has changed.

Typing something we do not hold

You are not restricted to the catalogue. If the server you have genuinely is not in it, type it and confirm — the line is saved exactly as you entered it.

The box will tell you, underneath, that the value is not in the supplier catalogue and what the nearest entry is. Worth reading rather than dismissing: a `PowerEdge R650` and a `PowerEdge R650XS` are different machines, and if you meant the one we hold, the efficiency figures for the one we do not will be a rating requested by name rather than a catalogue match.

What the row colours mean

Green	Your value and ours agree
Amber	They differ — either we corrected your input, or you have changed ours

They compare the two sides of the row, nothing else. Picking a different catalogue entry therefore turns a row amber rather than green, which is correct: it is telling you the line no longer says what your file said. Whether the value is in the catalogue is the separate note underneath the box.

A CPU is compared on its model, not on how it was written. The clock speed, `(R)` and `(TM)`, and the words "CPU" and "Processor" are ignored, so `Intel(R) Xeon(R) Gold 6254 CPU @ 3.10GHz` is green against `Intel Xeon Gold 6254`. A different suffix is a different chip: a `6338` against a `6338N` stays amber.

You can also **skip** a line — useful when the server genuinely is not in the catalogue. Skipped lines are left out of the analysis entirely rather than guessed at.

> **Confirm the flagged lines before you run**
>
> An unconfirmed match still runs, but the result is only as good as the hardware it was matched to. Flagged lines are flagged because we could not tell — a minute spent here is worth more than any other check.

Editing lines

Any line can be edited after import: vendor, model, CPU, CPU count, memory, quantity and reference.

A matched line opens on its match. The server, the CPU and (where there is one) the GPU are shown separately. For each you see:

- **In your file**: what your file said.
- **Matched to**: the catalogue entry it was matched to. It is shown in red if the catalogue does not hold it.
- **How it was decided**: *matched automatically*, *confirmed by hand*, *awaiting review*, or *carried from line N* when a decision made on another line in the same import was applied to this one.

This makes it easy to check someone else's matching. Saving without changing anything leaves the match exactly as it is, so correcting a quantity or a memory figure never undoes a match you confirmed.

To change one, press **Change server**, **Change CPU** or **Change GPU**. Only that part opens, as a box that lists what the matcher considered for your file's text, with the current value marked *current*, and searches the catalogue as you type. The other parts are left alone. If you enter something the catalogue does not hold, the box says so and suggests the nearest entry. You can still save it, but it will bring no core count or grade. Whatever you change, your file's original text is kept, and **Show original input** still shows it.

A line that did not match opens with what your file said, ready to correct. When you save, it is matched against the catalogue again.

Lines can be deleted individually, or several at once.

The server list on an EU compliance report behaves the same way, and adds the drive counts and utilisation that the label needs.

Reading the server list

Every line shows one status. There are two questions: is its server and CPU in the supplier catalogue, and has it been graded.

Status	What it means	What to do
Graded	In the catalogue, and graded	Nothing
Grading…	In the catalogue; the grade is on its way	Nothing. Starting a run fetches anything outstanding first
Not graded	In the catalogue, but our supplier has no grade for it. Rare	Change the server or CPU to one that is graded, or delete the line
Not in catalogue	Did not match, or was matched to a server or CPU the catalogue does not hold. The one it does not hold is shown in red	Edit the line and pick a catalogue entry, or delete it
To review	The match is waiting for you in match review	Open **Review Matches**
Skipped	Skipped during match review, and left out of the analysis	Nothing, unless you want it back

The line above the list counts each status. The filter buttons show one status at a time. To clear out the lines you cannot use, filter to one status, tick the box at the top of the list to select every line in it, and delete. The model and CPU columns show the catalogue's names; **Show original input** switches them to what your file said.

The core total above the list only counts lines whose CPU has a core count from the catalogue. When some do not, it says how many lines were left out.

Retry matching appears only when some lines failed to match or never finished, for example because the service was unavailable during import. It runs matching again for those lines. Lines that already matched are left as they are.

What if the hardware is not in the catalogue?

The catalogue covers a large range of server configurations and is refreshed daily. If something genuinely is not there, get in touch — see Getting help.

Importing from CSV

The import is a four-step wizard: **upload**, **map columns**, **preview**, then **confirm**. Nothing is imported until the final step, so you can back out at any point.

1. Upload

Choose **Upload specification** on the simulation page, pick your file, and choose **Upload**.

Your file is read but nothing is imported yet.

> **Use the template if you are starting from scratch**
>
> A template CSV is linked on the upload screen, already in the shape we expect. Exporting from your existing inventory tool and mapping the columns is usually faster than filling the template in by hand.

2. Map your columns

Every row on this screen is a field we understand. The dropdown beside it is the column from your file that we think matches.

We detect vendor, model, CPU and memory fields automatically where we can. The summary tells you how many we matched and how many still need you.

Fields marked **Required** in red must be mapped before you can continue. Everything else is optional and only adds detail to the results — but the more you map, the more precise the analysis.

Change any dropdown that looks wrong.

3. Preview

Your first few rows, arranged the way we have understood them.

This is the moment to spot a column that has landed in the wrong place. If these rows read correctly, the rest will too.

4. Confirm

Choose **Import**. The rows come in and each is matched against the hardware catalogue.

You are returned to the simulation, where anything we were unsure about is flagged for you to confirm. See Specifications.

File requirements

- Format: CSV.
- A header row naming the columns.
- One row per server configuration.
- Uploads are limited to 10 MB. If your file is larger, split it and import in parts — the lines accumulate on the specification.

Common problems

What you see	Usually means
Columns detected but all unmapped	The header row is missing, or is not the first row
Every row flagged after import	Vendor and model landed in the wrong columns — check the preview
Rows silently missing	Blank lines or a trailing separator row in the file
Upload refused	The file is over 10 MB, or is not a CSV

Targets and runs

A **target** is hardware you are considering. A **run** is one execution of the analysis, comparing your specification against every target you have added.

Adding targets

From the simulation page, choose **Add Target** and enter the vendor, model and CPU model. Suggestions come from the catalogue as you type. Then set the CPU count and memory.

Each target shows a status, the same one the server list uses: **Graded**, **Grading…**, **Not graded** or **Not in catalogue**. A part the catalogue does not hold is shown in red. See Reading the server list for what each status means. A target is usually graded straight away if we have graded that configuration before, otherwise within a minute or so of adding it.

A graded target can also say **No core count**, when the catalogue has no core count for its CPU. A run still goes ahead, but that target's licensing costs are estimated from your current estate's cores per server.

Editing a target shows its server, CPU and (if it has one) GPU separately, each as the catalogue entry it is matched to. For a target added from a SKU, it also says which SKU. Saving without changing them keeps them, so correcting the memory or the number of servers never changes the hardware. **Change server**, **Change CPU** or **Change GPU** opens that part only, as a box that searches the catalogue as you type.

Refresh Ratings asks for any grades that are missing or out of date. It queues the request rather than waiting for it, so the grades appear on the page shortly afterwards rather than the moment the button returns.

It cannot help a line marked **Not graded** or **Not in catalogue**. For a line that is not graded, our supplier has already answered that it holds no grade for it, and asking again gets the same answer. For a line that is not in the catalogue, the supplier does not know that server or CPU at all. In both cases, change the line to the nearest catalogue entry, or get in touch if the hardware is right and should be in the catalogue. See Reading the server list.

If your organisation has saved its standard builds, **From SKU** beside **Add Target** lists them with their status: tick any number (on a host-to-host simulation, each with its own quantity, 1 if you leave it empty), and each specification is copied in, already matched against the catalogue. Your SKU reference is carried onto the target and shown in the results and the PDF report. See Server SKUs.

Add as many targets as you like. Each one is analysed independently and they appear side by side in the results.

> **Each target costs one credit per run**
>
> Five targets means a run costs five credits, every time you run it. See Credits and billing.

Simulation types

Choose the type on the simulation page or at submission. You can change it at any time, as long as a run is not in progress.

Type	What it does
Estate Refresh	Compares your whole estate against each target and works out how many of that target you would need to replace it. The default.
Host to Host	You state exactly how many target servers you intend to buy, and the two fleets are compared directly.
Discovery	Grades the estate you already run, with no target and no comparison. A flat one credit.

Estate Refresh answers "if I replaced everything with this, what happens?". Host to Host answers "I am buying twelve of these — what changes?". Discovery answers "what have I actually got?" and is the right starting point when you are not yet comparing anything.

Your organisation may also have GPU variants of Estate Refresh and Host to Host available, which consider GPU-bearing servers only. If you do not see them, they are not enabled for your organisation.

> **Switching back to Estate Refresh clears target quantities**
>
> Host to Host needs a quantity per target; Estate Refresh calculates them. Switching back therefore blanks the quantities you entered.

Statuses

Status	Meaning	What you can do
Draft	Created, no specification yet	Upload a CSV, add lines, import from a location
Validating	Lines are being matched, or are waiting on your review	Review matches, add or edit lines, add targets
Ready	No line waiting for review, at least one target added. Unmatched lines are left out of the run	Submit, add more targets, edit lines
Processing	A run is in progress	Wait. Editing is blocked
Paused	A run is waiting for API capacity and will continue on its own	Wait — see Paused runs. Do not resubmit
Complete	At least one run has finished	Read results, generate reports, submit further runs

Cost parameters

Setting cost parameters is what unlocks the operational cost figures in your results and reports. Three are available:

Parameter	What it is	Example
Maintenance cost per server	Annual cost of keeping one server running — support contract, parts, staff time	1,200.00
Licensing cost per core	Annual per-core licensing for software priced that way	50.00
Energy cost per kWh	What you actually pay for electricity	0.15

Set them on the simulation page and they apply to every future run. You can also adjust them on the confirmation screen for one run only.

Anything you leave blank falls back to our platform assumption, and the report states plainly which figures were assumed rather than supplied. Supplying your own energy price is the single change that most improves the accuracy of the cost result.

Submitting a run

Choose **Run Simulation**. The confirmation screen shows how many source servers and targets will be analysed, what it costs, and what your balance is. Nothing is charged until you confirm.

A run usually takes a few minutes. You do not need to wait on the page — the simulation updates itself when the result arrives.

If a run is still processing after an hour, it is cancelled automatically and the credits are returned to your balance. You do not need to do anything to recover them.

Paused runs

Occasionally a run shows **Paused** rather than Processing, with a note that it is waiting for API capacity. Our hardware efficiency provider limits how many requests we may make in an hour, and a large run can reach that limit part-way through.

Nothing has gone wrong, and there is nothing for you to do:

- The targets already analysed are kept — the run picks up where it stopped.
- Your credits stay with the run. Nothing is refunded, because nothing failed.
- It restarts by itself, usually within a few minutes.
- You can close the page. We email you when the results are ready, as with any other run.

Do not submit the run again — a second copy would wait behind the same limit and cost you a second set of credits.

When the run finishes it looks like any other completed run, because it is: a pause changes nothing about the results.

In the rare case that capacity does not free up within 24 hours, the run finishes with whatever it managed to analyse, and the credits for the targets it could not reach are returned to your balance.

Run history

Every submission creates a new run with its own version number. The run history on the simulation page lists them all, with status, version, when it was submitted and a summary of the result.

Older runs stay readable even after you change the specification or targets, so you can compare what different configurations produced. This is what makes the duplicate-and-adjust approach below worth using.

Duplicating for what-if work

Duplicate creates a fresh draft from an existing simulation. It carries over the name, customer details and notes, but starts with no specification and no targets.

The usual pattern is to build a baseline, duplicate it, swap in different targets, and compare the results.

Finding a simulation again

The **Simulations** page lists everything you can see: your own work, private and shared alike, together with the simulations your colleagues have set to **Shared**. Anything a colleague has left private is not there — see Your account. The newest are at the top.

Three controls narrow the list, and they combine:

- **Search** matches on the simulation name, the customer name and the customer reference. Filling in customer details when you create a simulation is what makes this useful six months later — a reference number finds the job when nobody remembers what the simulation was called.
- **Status** shows only simulations at one point in their lifecycle — Draft, Validating, Ready, Processing or Complete, the same five described under Statuses. **Validating** picks up whatever was left mid-review, and **Ready** lists everything waiting to be run.
- **Folders**, in the sidebar, filter to one folder. **No Folder** gathers everything unfiled, which is the quickest way to find work that was never put away.

If a simulation you expect is missing, a filter is the usual reason — clearing the search box is not enough on its own when a folder is also selected.

Folders

Folders group simulations by project, client, or however you prefer. Create one from the sidebar on the Simulations page, and assign a simulation to it when creating or editing. Each folder shows how many simulations it holds.

Reading your results

A finished run answers one question in five dimensions: what changes if you move from the estate you have to the hardware you are considering.

The five measures

Energy reduction

Annual kWh saved, and the same figure as a percentage. This is the primary result — everything else follows from it.

Carbon reduction

Scope 2 emissions avoided, in CO_2e, derived from the energy reduction and the carbon intensity of your grid.

Server consolidation

How many servers the target hardware needs to do the same work, and the floor space that releases. Modern hardware is usually denser, so the count falls.

Energy rating improvement

Efficiency grades for your current estate and for each target. This is the fastest way to see whether a given piece of hardware is worth considering at all.

Cost savings

Annual operational cost reduction, built from the energy reduction and whatever cost parameters were in force for the run.

Servers left out of a run

A run analyses the lines of your specification that are matched to our catalogue. Lines that are still unmatched when you submit are left out, and every figure (servers, cores, memory, energy, the estate's rating) covers the matched servers only. When that happens, the results page and the report say how many specifications and servers were left out. To include them, match the lines and run the simulation again.

What the colours mean

Every figure that shows a change is coloured by whether the change is good for you: **green** when it is better, **red** when it is worse, and uncoloured when nothing changed. Better depends on the figure. Less energy, less carbon, fewer servers, fewer cores, a lower running

cost and a lower power draw are better; more performance is better; a change in memory is neither, so it is not coloured.

That means a negative reduction is red. A carbon reduction of −7% is a carbon increase, and shows red. The sign and the wording ("Increase", "Cost Increase") always stay beside the colour, so a report printed in black and white still reads correctly. The results page and the PDF report always colour a figure the same way.

Power and energy

Each target also shows its power and energy against your current estate:

- **Annual energy**, in kWh a year, including cooling and power distribution.
- **Power draw**, in watts at peak and at idle, for all the servers together: every server in your estate today, and as many of the target as the simulation says you would need. This is the servers alone.

Power draw needs a wattage for every validated server and for the target. If any is missing, that figure shows as not available rather than as a partial total. GPU simulations do not show power draw. These figures are worked out when a run finishes, so a run completed before they were added does not show them; run the simulation again to see them.

Efficiency grades

Every configuration in our catalogue is graded from **A+** down to **F**, based on its ITEE score — a measure of how much useful compute a server returns for the electricity it draws. Higher is better.

Grade	ITEE score	What it means
A+	14,000 and above	Substantially above average. Current, high-performance hardware
A	12,000 – 14,000	Well above average. Recent generation
B	9,000 – 12,000	Above average. Reasonably current
C	7,200 – 9,000	Moderate. Likely a few generations behind
D	6,000 – 7,200	Below average. Ageing, worth reviewing
E	4,700 – 6,000	Poor. A strong candidate for replacement
F	Below 4,700	Very low. Replace these first
U	—	Ungraded. We could not determine a rating

Grades appear as coloured badges throughout: on specification lines, on targets, on server detail pages, and in the dashboard charts. They are the quickest way to see the shape of an estate and where the opportunity sits.

Where the grades come from

Grades are not our opinion. Each configuration is rated against industry benchmark data — peak power draw, idle power draw, and benchmark throughput — and the ITEE score falls out of that. Ratings are fetched when a configuration first appears, and the catalogue is refreshed daily so new hardware and revised ratings arrive without anyone asking.

What a U grade means

U means we could not rate that configuration. Usually one of:

- The hardware is new enough that it is not yet in the ratings data.
- The vendor, model or CPU does not match anything we know — most often a typo or a non-standard name in the source inventory.
- The rating service was briefly unavailable when we looked.

The fix is nearly always the second one. Check the details on the line, correct anything wrong, and it will be rated on the next attempt. See Specifications.

Ungraded lines still appear in your specification, but contribute nothing to the energy calculation — so a simulation with many U grades understates both the current estate and the saving.

How far to trust it

The analysis is only as good as three things:

1. **How well your estate was matched.** Lines flagged and left unconfirmed, and lines left at U, are the main source of error. See Specifications.
2. **Whether you supplied your own energy price.** The default is a platform assumption and will not be your tariff.
3. **How complete your import was.** Optional fields you did not map still sharpen the result.

The report states which figures were assumed rather than supplied. Read that section before circulating the result.

Comparing targets

Where a simulation has several targets, the results place them side by side. The best choice is rarely the one with the largest headline energy saving — consolidation and licensing often move the cost result further than energy does.

When part of a run fails

If some targets fail and others succeed, the run still completes. The failed targets carry their error message, and their credits are refunded automatically. You are never charged for analysis that produced nothing.

Exporting

Every result can be exported as a PDF report. See Reports.

EU COMPLIANCE

EU compliance reports

An EU compliance report states your data centre's efficiency in the terms the Energy Efficiency Directive uses: a **PUE** and a **WUE** figure, each placed on an A–G scale. It is a separate thing from a simulation — there are no targets and nothing is compared. You are describing one data centre as it stands.

Compliance is a module. If you do not see it in the navigation, your organisation does not have it enabled; ask your administrator or get in touch.

Building a report

1. **Create the report.** Name it, optionally record the customer and reference it is being prepared for — these are searchable from the report list and appear on the PDF cover, as they do for simulations — and set its location. The country list and the regions within each country come from our catalogue, so pick from the list rather than typing — the region has to be a code the compliance service recognises.
2. **Add your servers.** Either import them from a CSV or add lines by hand. This works exactly as a simulation specification does, including the matching and the grades, so anything you already know about that screen applies here. After a CSV import you land on a progress screen while the file is matched against the catalogue — a large estate takes a few minutes — and then, exactly as with a simulation, any uncertain matches are brought to the **review screen**, where each one can be confirmed, corrected to an alternative, or skipped. The report cannot be run until the review is finished. If matching fails, the reason appears on the report and you can retry by editing a line or importing again; the imported servers are never lost.
3. **Enter your facility figures.** PUE and WUE are yours to supply. We do not measure them. Set **Server utilisation (%)** here too — it drives ITEU and the average power that the energy, carbon and water figures are all built from.

   Networking share (%) and **Storage share (%)** are the parts of your IT load drawn by networking and dedicated storage equipment. They raise IT energy above the servers' own consumption, so leave them at 0 only if you genuinely don't measure them separately. Together they must be under 100% — the rest is the servers.

   Utilisation is taken from the best information available, server by server: a figure imported against a server is used as it stands; where a server has none, this facility average stands in; and where neither exists, that server is calculated at 25%. Importing a `utilisation` column for the servers you have measured and setting this average for the rest gives the most accurate result.

   Accelerators work the same way, through their own `gpu_utilisation` column and **GPU utilisation (%)** figure. 4. **Run it.** One credit, however many servers are in the estate. See Credits and billing.

A run takes a few minutes. The page follows it and takes you to the results when it finishes; you do not have to wait on it.

Reading the results

Each metric is shown three ways: the measured value, the grade it earns, and its position on the A–G scale, with the two scales facing each other as they do on the printed label. The bands and thresholds are the scheme's, not ours.

> **PUE and WUE are not validated**
>
> Both figures are the ones you entered. We place them on the scale and grade them, but we do not verify them, and the facility and total energy figures shown alongside are estimates derived from the PUE you gave us. The report is only as good as those two numbers.

Server analysis

Below the grades, every server line is listed with its own efficiency rating — the same A–F rating a simulation shows — alongside the average power and annual energy the compliance service calculated for it.

Refresh Ratings asks for any grades that are missing or out of date. It queues the request rather than waiting for it, so reload the report a minute or so later to see them.

The last column says what the catalogue made of each line. **Server not in catalogue** or **CPU not in catalogue** means that line could not be rated, and so contributed nothing to ITEE, power or water. Those lines still count toward your totals; they are listed rather than hidden precisely because they are the ones worth correcting. Fixing a match on the report's server list and running again will bring them in — see Specifications.

Every run is kept. The **Run History** panel lists each attempt, and opening one shows that run's own results — useful when a figure has been corrected and you want to see what changed.

The label

View Label renders the results in the layout of the EU data centre label — the identity block, the two grade scales, and the voluntary sustainability features.

Some boxes on it are empty, shown as **—**. Those are fields the label defines that the platform does not yet collect: the type and size category of the data centre, the reporting period, the date it entered operation, the renewable energy breakdown, waste heat reuse, and grid services. They are left blank on purpose. We would rather show you the gap than fill a regulatory document with a plausible guess, and the list is what we are working through.

The PDF

Download PDF produces the report as a document. **Customise** controls what goes in it — which sections appear, whether your own branding is applied, and the cover and summary text — in the same way as simulation reports.

REPORTS AND CREDITS

Reports

A report turns a finished run into a document you can send to somebody who has never opened Interact DC.

Generating one

Open a finished run and choose **Generate Report**. The report is built in the background and downloads as a PDF when it is ready.

> **Quick Generate skips the configuration**
>
> If the defaults suit, **Quick Generate** produces the report in one step.

Choosing what goes in

The configuration screen lets you include or leave out each section:

Section	What it holds
Cover page	Customer details from the simulation, and your branding
Executive summary	The headline figures, for a reader who will go no further
Source analysis	Your current estate, summarised — counts, grades, energy
Target comparisons	Each target side by side against the source
Methodology	How the figures were arrived at, and which were assumed
About	Who Interact DC are

Individual metrics can be toggled on or off within those sections, and you can add your own notes. Under target comparisons, **Annual Energy** (kWh a year, current against new) and **Power Draw** (peak and idle watts, current against new) are both ticked by default; untick either to leave it out.

That methodology section matters. A reader who does not know which numbers came from you and which came from our defaults cannot judge the result, so it is stated plainly rather than buried. Think twice before removing it.

Branding

If your organisation has configured branding — logo, colours, company name, contact details, disclaimer — reports use it automatically. You can toggle those elements per report. Without it, reports carry Interact DC branding.

An organisation administrator sets it from the branding page. See Administering your organisation.

Before you send it

- Confirm any flagged specification lines and re-run, if you have not already.
- Supply your own energy price rather than accepting the default.
- Fill in the customer details on the simulation — they appear on the cover.
- Check for **U** grades in the source analysis. Many of them means the estate is under-represented.

Sharing

Set the simulation's visibility to **Shared** so colleagues in your organisation can open it themselves rather than working from a forwarded PDF. See Your account.

Credits and billing

Running a simulation costs **credits**. Credits belong to your organisation, and every run is charged against that balance.

What a run costs

One credit per target, per run. Three targets cost three credits every time you run them. A **Discovery** run, which has no targets, is a flat one credit.

An **EU compliance report** is a flat one credit per run, however many servers are in the estate. A 40,000-server data centre and a 10-server one cost the same, because it is one report either way.

You are always shown the exact figure on the confirmation screen **before** anything is charged — there is no way to spend credits by accident.

Building is free. Creating a simulation, uploading your estate, mapping its columns and choosing targets all cost nothing, and you can do all of it with a zero balance. Credits are spent at the moment you run. Where a run is not affordable the button is disabled and says what it would cost, so you can see the price before you commit to the work.

Seeing your balance

Your organisation's balance is shown in the app. Organisation admins can see the full history of what was bought and what was spent.

Personal allocations

Your organisation may give you a personal allocation — credits set aside for you, carved out of the shared pool.

When you submit a run, your personal credits are used first. If they run out and you are permitted to draw on the shared pool, the rest comes from there. If you are not, the run is refused until your allocation is topped up.

Administrators set both the allocation and that permission — see Administering your organisation.

Somebody whose role is **Viewer** cannot be given an allocation. A viewer has no way to spend credits, so the allocation would leave the shared pool and sit there unusable; change their role first if they need to run simulations.

Refunds

If a run fails, the credits come back automatically. If every target failed you get the lot back; if some succeeded and others failed, only the failed ones are refunded.

A run that is still processing after an hour is cancelled and refunded on the same terms.

You are never charged for analysis that produced no result.

Refunding a purchase

If you need the money back for credits you have bought and not used, contact support and we will arrange it. Refunds can cover part of a purchase — you do not have to reverse the whole thing.

Unused credits are refundable. Credits you have already spent on a simulation are not, because the analysis and its report have been delivered.

When a refund is made, the credits are removed from your balance, the money is returned to the payment method you used, and a credit note appears on your credits page alongside the original invoice.

Buying more

Organisation admins can buy credits from the credits page, if self-service purchasing is enabled for your organisation. Payment is taken by card or PayPal, and a receipt is issued for every purchase.

Credits are added to the balance as soon as the payment clears.

> **Confirm your email address first**
>
> Your receipt is a VAT document and is emailed to you, so we ask you to follow the confirmation link we sent when you signed up before you can buy. If you have lost it, opening the buy-credits page sends a fresh one. Nothing else needs it — building a simulation does not.

If self-service purchasing is not enabled, your administrator arranges credits directly with us.

Receipts and credit notes

Every purchase produces a numbered receipt, available to download from the credits page. Admins can also export the full receipt history.

A refund produces a **credit note** rather than cancelling the original receipt. The receipt stays as the record of what you bought; the credit note is the record of what was returned, and the two are read together. Both are numbered, both are downloadable, and both are included in an export.

Running out mid-project

A run that would take you below zero is refused at the confirmation screen rather than started and abandoned, so you never lose a partially-charged run. Top up and submit again — nothing about the simulation is lost.

YOUR ORGANISATION

Administering your organisation

If you are an administrator, a set of **Organisation** pages appears in the left sidebar that other members do not see: **Dashboard**, **Members**, **Credits** and **Analytics**. Settings, branding and the full activity log are reached from the dashboard.

Everything on these pages is scoped to your own organisation. Nobody outside it can see any of it, and you cannot see anyone else's.

> **If you work on your own**
>
> An account that signed up by itself sits in a personal organisation with one member. There is nobody to administer, so the pages below do not appear — but you still hold your own credits, and the credits page is still yours.

Dashboard

The opening page summarises the organisation:

- Credit balance, and whether it has fallen below your alert threshold.
- Member count, and how many accounts are suspended.
- Simulations in total and this month, and runs over the last 30 days.
- The most recent activity and credit transactions.
- Six months of activity as a chart.

Branding and **Settings** are the two buttons at the top of this page.

Members

Members is where people are added, changed and removed. It lists everyone in the organisation with their role, status and personal credit allocation, and any pending requests to join above the table.

Getting somebody an account

Invite User sends an invitation to an email address. The person receives a link, chooses a password, and lands inside your organisation.

An invitation goes out as a **Member** unless you change the role before sending it. Members can do the work; only an administrator can invite people, manage roles and spend credits, so that is worth choosing deliberately rather than by default.

No password is ever emailed. If an invitation goes unanswered or expires, use **Resend Invitation** rather than creating a second account — an account showing **Invited** exists already and is simply waiting for its owner.

Somebody who already has an account can ask to join your organisation instead. Pending requests appear at the top of the members page, and **Approve** or **Decline** settles them. Either way, they are emailed the outcome.

Roles

Role	Can do
Viewer	Read work shared with the organisation and download or generate reports from it. Cannot create, edit, run anything or spend credits
Member	Everything a viewer can, plus create and run simulations and spend credits
Administrator	Everything a member can, plus everything on this page

Change somebody's role with the edit action on their row.

Viewer is the right role for somebody who needs the results and should never be able to spend the organisation's credits — an executive sponsor, a sustainability lead, an auditor. They see what has been shared with the organisation, not private work.

Two things a viewer cannot be given: a personal credit allocation, since they can never spend it, and ownership of the organisation, which only an administrator can hold.

A viewer sees the same pages as everyone else. Controls they cannot use are greyed out rather than removed, and hovering one says why — so a viewer asking "where is the New Simulation button?" can see that it is there and why it is unavailable, rather than wondering whether something is broken.

One administrator is also the **owner**. The owner cannot leave the organisation while they own it — use **Make owner** on another administrator's row to hand it over first. There is always exactly one owner.

Suspending and removing

Suspend blocks sign-in immediately and leaves everything else intact. Their simulations stay where they are, and reactivating restores access. This is the right action for somebody who has left, or an account you are unsure about.

Delete User removes the account. Any simulations they own are transferred to you rather than deleted, so nothing is lost — but the account itself is gone and cannot be restored. Suspension is reversible; deletion is not.

Reset Password emails them a link to set a new one. You never see or set a password on somebody's behalf.

Credits

Credits holds the balance, the full transaction history, and your VAT receipts for anything purchased.

The organisation has one shared pool. From the members page you can carve a **personal allocation** out of that pool for an individual: credits set aside for them, deducted from the shared balance when you allocate them.

Each allocation also carries **Allow shared pool access**. With it on, that person's runs draw on their allocation first and fall back to the pool when it is empty. With it off, they can spend only what you have given them and their runs are refused once it is gone — which is how you cap somebody's spending.

The transaction history records every allocation, purchase, consumption and refund, with who performed it. Failed runs refund themselves; see Credits and billing.

Settings

From **Settings** on the dashboard:

Setting	What it does
Name	The organisation name, as it appears in the app and on reports
Default simulation visibility	Whether new simulations start private or shared
Credit alert threshold	The balance at which low-credit warnings are sent to administrators

The alert threshold is what drives the low-credit email. Administrators can turn that email off individually — see Email notifications.

Branding

Branding on the dashboard is where you set what appears on generated PDF reports: company name, logo, primary and secondary colours, contact email and phone, website, an About paragraph, and a report disclaimer. **Powered by Interact DC** can also be turned off here.

Reports use this automatically once it is set — there is nothing to switch on per report, although individual elements can be toggled off when generating one. Without branding, reports carry Interact DC's own.

A logo must be a PNG, JPEG, GIF or WebP of no more than 2 MB. A wide, transparent PNG reproduces best on a report cover.

See Reports.

Server SKUs

Server SKUs on the dashboard is where you save your standard server builds against your own references, so that members add them as simulation targets by picking from a list rather

than retyping a specification. Builds can be typed in one at a time or imported from a spreadsheet.

See Server SKUs.

Activity log

The dashboard's recent activity panel links to the full log. It records who did what and when — sign-ins, member changes, credit movements, settings and branding changes — and filters by activity type and by person.

This is the page to open when you need to answer "who changed this, and when".

Analytics

Analytics covers a period you choose — 7 days, 30 days, 90 days or a year — and reports simulations created, runs submitted, how many completed and how many failed, credits consumed, and who has been most active.

A run count well above the credits consumed usually means work is being re-submitted; a high failure count is worth telling us about.

Server SKUs

A **server SKU** is one of your standard builds, saved against your own reference — `OS-32C-BASE-R7715`, `GPU-8xH100`, whatever you already call it internally. Save it once, and anyone in your organisation adds it as a simulation target by picking it from a list instead of retyping the specification.

Your reference then travels with the work: it appears on the target, in the run results, and in the PDF report.

> **Administrators define them, everybody uses them**
>
> The SKU list is the vocabulary your organisation's simulations are described in, so defining and editing it is an administrator's job. Every member can pick from the list.

The SKU list

Organisation Admin → Server SKUs lists every build you have defined, with its status. We match each SKU against the hardware catalogue when you save it, so the list tells you straight away whether a build is ready to use.

The statuses are the same as on a simulation's server list:

Status	What it means	What to do
Graded	In the catalogue, and graded	Nothing
Grading…	In the catalogue; the grade is on its way. The page updates itself, usually within a minute	Nothing
Not graded	In the catalogue, but our supplier has no grade for it	Change the server or CPU to one that is graded
Not in catalogue	The build did not match, or part of it is not in the catalogue. That part is shown in red	Edit the SKU and pick a catalogue entry
To check	We matched the build, but only as a best guess	Edit the SKU, check the match, then **Confirm match** or change it

The grade itself appears in your simulation results, not here.

The model and CPU columns show the catalogue's names for what the SKU matched. **Show your definition** switches them to what you typed. The line above the list counts each status, and the buttons under it show one status at a time.

The **Cores** column comes from the configuration we matched, not from what was typed — so it is also a check on the match. A build called `OS-32C` whose cores read 16 has matched the wrong CPU.

Defining one

New SKU asks for:

Field	Notes
Reference / SKU	Your own code. Must be unique within your organisation
Vendor, **Model**, **CPU**	What the build is, as the catalogue would name it
CPUs, **Memory (GB)**	The configuration
GPU, **GPU count**	Optional. Leave blank for a CPU-only server. Only shown if GPU support is enabled for your organisation

Vendor, model, CPU and GPU suggest as you type, from the same catalogue the **Add Target** form uses, and a build we cannot place exactly shows a "did you mean" panel underneath with the closest matches. Taking the suggestion is worth it: a SKU is typed once and used for years, so an unrecognised spelling here costs a grade on every simulation that picks it.

There is no quantity. A SKU says what a server is, never how many you own — how many you would buy is asked when the SKU is used, because that is a question for one simulation rather than for the definition.

Importing a spreadsheet

Most organisations already keep their standard builds in a spreadsheet. The dropdown beside **New SKU** has **Import from CSV**, which takes it as it is.

1. **Upload** your file. Nothing is imported yet.
2. **Map your columns.** Every row is a field we understand, and the dropdown beside it is the column from your file we think matches. We recognise the usual spellings — `Make`, `Sockets`, `RAM (GB)`, `SKU Reference` — so the screen usually comes up answered. Change anything that looks wrong. **Reference**, **vendor** and **model** must be mapped.
3. **Import.** The SKUs are created and matched against the catalogue, and the list shows each one's status. Filter to **Not in catalogue** or **To check** to see the ones that need attention.

Any other columns are ignored, including a quantity column — see above.

Download Template in the same dropdown gives you a CSV in the shape we expect, if you would rather start from scratch.

Rows that are refused

A reference must be unique within your organisation, so a row is refused if its reference is already in use — by an existing SKU, or by an earlier row in the same file — or if it has no reference, vendor or model at all.

Refused rows are named individually, with the row number and the reason, above the list. Everything else in the file is imported: one bad row does not cost you the upload.

Using a SKU

On any simulation, **Add Target** has a **From SKU** option. It lists your organisation's SKUs with their status. Tick as many as you want, then **Add selected** adds a target for each.

On a host-to-host simulation, each SKU also has a quantity: how many of that build you would deploy. Typing a quantity ticks the SKU for you, and a ticked SKU with no quantity counts as one. An estate refresh works out the number of servers itself, so it does not ask.

A SKU that is already a target on this simulation is not added twice. If you give it a quantity, that target's quantity is changed to it; otherwise it is left as it is. The message says which.

Two things happen that are worth knowing about.

The specification is copied. The target takes the SKU's vendor, model, CPU, memory and GPU as its own, along with the SKU reference. It does not point back at the SKU. For a matched SKU, the vendor, model and CPU are the catalogue entries it is matched to — what the SKU table and the picker show — not what you typed. So a SKU whose server you have changed adds a target that is the new server.

The catalogue match is copied too, which is why adding a target this way is immediate: the matching was done when the SKU was saved.

Editing, retiring and deleting

Editing a matched SKU shows its server, CPU and (if it has one) GPU separately. For each you see what you typed (**In your definition**), what it is matched to, and how that was decided. Saving without changing them keeps the match, so correcting the memory or CPU count never undoes it. **Change server**, **Change CPU** or **Change GPU** opens that part only, as a box that lists what the matcher considered and searches the catalogue as you type. Your definition is kept whatever you change. A SKU marked **To check** also has **Confirm match**, which accepts the match as it stands.

Editing a SKU that did not match shows its definition to correct. It is matched again when you save.

Editing does **not** change any target already built from a SKU. Those carry their own copy, so no finished simulation or report can be rewritten by a later edit. To bring an existing target in line with a corrected SKU, either edit the target and **Change server**, or remove it and add the SKU again. Rerun the simulation for the change to reach its results.

Retire takes a build out of circulation. It disappears from the picker, stays in the admin list marked **Retired**, and the targets and reports that already used it are untouched. Retiring can be undone.

Delete removes it altogether. Nothing downstream breaks: every target built from a SKU carries its own copy of the specification and the reference it ran with, so the simulations and reports you have already produced stay exactly as they are, still printing that reference.

Which to use:

Delete	A mistake, a test entry, a junk row from an import — anything nobody has used
Retire	A build you have moved on from but have run simulations against. It leaves the picker, and you keep the record of what you standardised on and what those simulations were comparing

The confirmation tells you how many targets were built from a SKU before you delete it, and suggests retiring instead when there are any.

Where your reference appears

- The target row on the simulation page.
- The target headings in run results.
- The target comparison table in the PDF report, where the column appears only if at least one target in the run came from a SKU.

Because the reference is recorded with the run, an old report keeps printing the reference it was run with, whatever has happened to the SKU since.

ESTATE

Locations and estate

Interact DC has two modules, and your organisation may have one or both.

The **Simulations** module is the planning tool: describe what you run, compare it against hardware you are considering, get a report. Everything else in this guide is about that.

The **Estate** module is the record of what you have today — locations, server inventory, utilisation over time, and dashboards over the lot. It is a live picture rather than a projection.

> **Estate is off unless you asked for it**
>
> New organisations start with Simulations only. If nothing in this page matches what you see, your organisation does not have the Estate module enabled — ask your administrator.

> **Changing the estate is an administrator's job**
>
> Everyone in your organisation can read the estate — the inventory, the dashboards, the utilisation history. Adding, editing or removing servers, locations and server groups is limited to administrators, so the record of what you own stays in one pair of hands.

The two work well together: keep the inventory current in Estate, then pull it straight into a simulation instead of re-entering it. Each also works on its own.

Locations

A **location** is a physical site — "London DC1", "Frankfurt Colo", "New York DR". Every server in the Estate module belongs to one, and locations are the main way of organising servers by where they actually sit.

Each location has a dashboard showing how many servers are active, their average efficiency grade, CPU and memory utilisation, and how those have moved over time.

The reason locations matter outside the Estate module: from a simulation you can choose **Import from Location** and pull that location's whole inventory into your specification in one step. See Run your first simulation.

Server groups

A **server group** is an optional grouping that cuts across locations. Where a location is about physical geography, a group is about whatever you care to organise by — project, business unit, application cluster, refresh priority.

A server can belong to several groups at once, and each group gets the same dashboard treatment as a location.

Bulk upload

The upload wizard brings server records or utilisation measurements in from CSV. It works in the same shape as the simulation import: upload, say what kind of data it is and which location it belongs to, confirm the column mapping, then process.

The same 10 MB limit and the same encoding and delimiter detection apply. See Importing from CSV.

Hardware configurations

Both modules are built on the same catalogue. A **hardware configuration** is one server specification — vendor, model, CPU model, CPU count and memory — and carries the efficiency grade that everything else is calculated from.

Servers sharing a specification point at one catalogue entry rather than each storing their own copy, which is why correcting a match once fixes it everywhere. See Reading your results for what the grades mean.

REFERENCE

Common tasks

A short index of the things people do most often, and where each one lives.

Task	Where
Create a simulation	**Simulations** → **New simulation**
Upload a server inventory	Open the simulation → **Source Specification** → **Upload CSV** → map columns → confirm
Add a line by hand	Open the simulation → **Add Data** → **Add Server Line**
Import servers from a location	Open the simulation → **Import from Location**
Confirm flagged matches	Open the simulation → **Review Matches**
Add a comparison target	Open the simulation → **Target Configurations** → **Add Target**
Set cost parameters	Open the simulation → cost parameters panel
Run a simulation	Open the simulation (status **Ready**) → **Run Simulation** → confirm
Find a past simulation	**Simulations** → search by name, customer name or reference
Read a past run	Open the simulation → run history → pick a version
Generate a PDF report	Open a finished run → **Generate Report**
Generate a report with defaults	Open a finished run → **Quick Generate**
Duplicate for a what-if	Open the simulation → **Duplicate**
Share with your team	Open the simulation → set visibility to **Shared**
File a simulation in a folder	Create the folder in the Simulations sidebar, then set it on the simulation
Change your password	Profile menu → account settings
Add your name to your profile	Profile menu → account settings
Check your credit balance	Credits page, or ask your organisation admin
Buy credits	Credits page, if enabled for your organisation
Invite a colleague	**Organisation** → **Members** → **Invite User** (admins only)
Approve a request to join	**Organisation** → **Members** → **Approve** (admins only)
Give somebody a personal allocation	**Organisation** → **Members** → allocate credits (admins only)
Set branding for reports	**Organisation** → **Dashboard** → **Branding** (admins only)
See who changed what	**Organisation** → **Dashboard** → recent activity → **View all** (admins only)
Choose which emails you receive	Profile menu → account settings → **Email Notifications**

Task	Where
Leave an organisation	Organisation page → leave

Tasks below need the Estate module, which is off unless your organisation asked for it. See Locations and estate.

Task	Where
Add a server to your estate	**Hardware** → **New Server**
Create a location	**Hardware** → **Locations** → **New Location**
Upload estate data in bulk	**Upload** → choose the file, type and location → map columns

Glossary

Credit
: The unit a simulation run is charged in. One credit per target per run; a Discovery run is a flat one. Held per organisation. See Credits and billing.

Discovery
: A simulation type that grades the estate you already run, without comparing it against anything.

Efficiency grade
: A grade from **A+** to **F** describing how efficiently a hardware configuration converts energy into work, derived from its ITEE score. **U** means we could not determine one. See Reading your results.

Estate Refresh
: The default simulation type. Compares your whole estate against each target and works out how many of that target would be needed to replace it.

Facility overhead
: How much energy the facility draws for every unit drawn by the IT equipment — cooling, distribution losses and the rest.

Hardware configuration
: One server specification in our catalogue — vendor, model, CPU model, CPU count and memory — carrying the efficiency grade everything else is calculated from.

Host to Host
: A simulation type where you state exactly how many target servers you intend to buy, and the two fleets are compared directly.

ITEE score
: IT Equipment Energy Efficiency. A measure of useful compute returned per unit of electricity drawn. Determines the letter grade.

Location
: A physical site in the Estate module. Every server there belongs to one, and a location's inventory can be imported straight into a simulation. See Locations and estate.

Match
: The catalogue hardware that one line of your specification has been identified as. Matching is fuzzy and carries a confidence score.

Organisation
: The account that owns simulations, estates, reports and credits. Every user belongs to one, and organisations are sealed from one another.

Run
: One execution of a simulation. A simulation keeps every run you have made, each with its own version number.

Scope 2 emissions

Emissions from the electricity you purchase. What the carbon result measures.

Server group
An optional grouping of servers that cuts across locations — by project, business unit, or however you prefer.

Simulation
The container for one piece of analysis: your estate, the hardware you are comparing it against, and every run.

Specification
Your current estate as described to Interact DC — one line per server configuration. See Specifications.

Specification line
One row of a specification, representing one or more identical servers.

Target
A hardware configuration you are considering, evaluated against your specification. Each costs one credit per run.

Visibility
Whether a simulation is *private* to you or *shared* with your organisation.

Getting help

Inside the app

A guided tour runs on the real screens and covers the whole path from an empty simulation list to a finished result. Open the profile menu in the top right and choose **Run your first simulation**.

Troubleshooting

I cannot sign in

Check you are using the email address your account was created with — there is no separate username.

If you have never signed in, you have no password yet: your welcome email carries a link to choose one. That link works once and expires, so if it has been sitting in your inbox a while, use **Forgot your password?** on the sign-in page instead — it gets you to the same place.

If your account has been suspended you will be told so; your organisation admin can reactivate it. Otherwise use **Forgot your password?**.

I cannot see a colleague's simulation

Simulations are private by default. Whoever created it needs to set its visibility to **Shared** before it appears for anyone else.

If it should already be shared, check you are not filtered to a folder — clear the filters on the Simulations page.

My CSV upload failed

The file must be a CSV under 10 MB with at least vendor and model columns. UTF-8 is safest. We detect commas, semicolons and tabs, so an unusual delimiter is rarely the cause — but if the mapping screen shows data in the wrong places, the header row is the first thing to check.

The template linked on the upload screen is already in the expected shape. See Importing from CSV.

My simulation is stuck processing

Runs usually finish in a few minutes. If one is still processing after an hour it is cancelled automatically and your credits are returned — you do not need to do anything to recover them.

If runs are consistently taking that long, tell us.

Some servers show a U grade

We could not rate that hardware, most often because the vendor, model or CPU does not match anything we know. Check those fields for typos. See Reading your results.

I do not have enough credits

Each target costs one credit per run, so five targets cost five credits. A Discovery run is a flat one credit.

If your organisation has self-service purchasing enabled, buy more from the credits page. Otherwise ask your organisation admin, who can allocate credits to you or top up the shared pool. See Credits and billing.

A run failed but I was not charged

That is deliberate. Credits for failed targets are refunded automatically — the whole lot if every target failed, or just the failed ones if some succeeded. You are never charged for analysis that produced nothing.

I get a 403, or a feature is missing

Not every module and feature is enabled for every organisation. The Estate module in particular is off unless your organisation asked for it. Ask your administrator.

Something failed with an error message

Error messages appear at the top of the page or beside the field concerned. The usual causes are a duplicate name, a required field left empty, or a value in an unexpected format. Read the message, correct it, and try again.

Reporting a wrong result

If a result looks wrong, check these first — they account for most of it:

1. Are any specification lines still flagged as unconfirmed? See Specifications.
2. Did the run use your energy price, or our default? The report says which.
3. Was the whole estate imported? Compare the line count against your source file.
4. Are many lines showing a **U** grade? Those contribute nothing to the energy figure, so the result will understate both sides.

If it still looks wrong after those, send us the simulation name and what you expected to see.

Contact us

Email support@interactdc.com.

To get a useful answer quickly, include:

- What you were doing, and what you expected instead.
- The name of the simulation, if it concerns one.
- Your organisation name.
- The address of the page, and anything the error message said.

Please do not send account passwords. We will never ask for one.

For account matters — password resets, role changes, credit allocations, suspensions — your organisation administrator can act faster than we can.

Terms, privacy and acceptable use

Four pages sit outside the app and need no sign-in: **Privacy**, **Terms**, **Cookies** and **Acceptable use**. They are linked from the footer and from the sign-in page.

Between them they cover what we hold, how long we hold it, what we do with it, and what the service may be used for. If you need a copy for a procurement or security review, or a question answered that they do not settle, ask us.

Something in this guide is wrong

Tell us. This guide is maintained alongside the application itself, so a correction can ship with the next release.

Feature requests are welcome too.