



TALA

User Manual

Version 0.2.13

Terrastruct

January 16, 2023

Contents

1	Introduction	3
2	Installation	3
3	Using TALA	4
4	License key	5
5	Considerations for software architecture	5
5.1	Symmetry	5
5.2	Clusters	6
5.3	Hierarchy	7
5.4	Balanced connections	8
5.5	Dynamic label positioning	9
5.6	Square aspect ratio	10
5.7	Direction	11
5.8	Near	13
5.9	SQL table column matching	13
5.10	Sequences	14
5.11	Self connections	14
6	Seed	15
7	Tips	16
7.1	Dimensions	16
7.2	Label collisions	17

8	FAQ	17
8.1	Will TALA require an internet connection?	17
8.2	Can I use TALA for diagrams which are not software architecture?	17
8.3	Can I use TALA to layout separately from D2?	17

1 Introduction

Thank you for purchasing a license for TALA.

TALA (short for Terrastruct’s AutoLayout Approach) is a proprietary layout engine designed for software architecture diagrams. The software is written entirely in Go, and was developed by Terrastruct over the span of 2 years. TALA is new, but already its current form can often produce diagrams that are preferred over the layouts generated by any other engine.

Unlike other layout engines, TALA is not one algorithm. It is comprised of many algorithms in multiple stages. It identifies parts of the diagram that are structurally like trees, or hierarchies, clusters, etc, and runs different algorithms for each. We used a combination of ideas from the decades of research done in the field of graph visualization, with modifications and new additions to tailor it to software architecture diagrams. More on this in [a following section](#).

If you find an issue in TALA (or this manual), please file it on [TALA’s GitHub](#)¹. If it’s a D2 issue, please file it on [D2’s GitHub](#)². Usually anything that pertains to the layout (sizing, positioning) falls under TALA. You can see how active we are in public and be sure that your issue will quickly be attended to.

We’d love to hear from you. Suggestions, ideas for integrating TALA into something, bugs, good things, bad things – all welcome. The above GitHub links are best, so that everyone can see/search and chime in on an open discussion, but we’re also around for live chat ~18hrs a day on [Discord](#)³. For private enquiries, email info@terrastruct.com.

2 Installation

The most convenient way to install is through the install script, together with D2.

If you want to upgrade versions, just re-run the script.

Alternatively, you can also find binaries [releases](#)⁴ for Linux and MacOS, for both AMD and ARM, and an installer for Windows.

¹<https://github.com/terrastruct/TALA>

²<https://github.com/terrastruct/d2>

³<https://discord.gg/aKXhwpCvrY>

⁴<https://github.com/terrastruct/TALA/releases>

```
# With --dry-run the install script will print the commands it will use  
# to install without actually installing so you know what it's going to do.  
curl -fsSL https://d2lang.com/install.sh | sh -s -- --tala --dry-run  
# If things look good, install for real.  
curl -fsSL https://d2lang.com/install.sh | sh -s -- --tala
```

Listing 1: Installing TALA

We also have a Homebrew tap, so you can install via Homebrew if you'd like (the install script uses Homebrew if detected on your machine).

```
brew install terrastruct/tap/tala
```

TALA's binary is named `d2plugin-tala`. To check that it was properly installed and in your PATH, run the following.

```
d2plugin-tala --version
```

To check that D2 can find TALA, run the following.

```
d2 layout tala
```

3 Using TALA

To use TALA for D2's layout engine, pass the flag `layout` flag:

```
d2 --layout tala input.d2
```

Shorthand:

```
d2 -l tala input.d2
```

You can also set it to an environment variable to not have to specify it each time:

```
export D2_LAYOUT = tala
d2 input.d2
```

4 License key

The easiest way to reference your license key is to put it in your PATH.

```
export TSTRUCT_TOKEN = "tstruct_..."
```

You can also write the key to a file. By default, TALA will look for this file in your home directory's `.config/tstruct/auth.json`. You can override that.

```
export TSTRUCT_AUTHFILE = "/home/my/path/config.json"
```

5 Considerations for software architecture

Instead of relying solely on theoretical cross-minimizing heuristics to dictate layout like other engines might, TALA emulates how software engineers might draw diagrams on a whiteboard. What follows is a non-exhaustive list of considerations that TALA makes towards that goal.

5.1 Symmetry

Symmetry is an important part of aesthetics and its weighted in TALA's algorithm. There are instances where a shape could've been closer to its neighbor, but doing so would've broken symmetry, and so it gets placed in the position that gives it more symmetry.




Figure 1: Example of symmetry

5.2 Clusters

Clusters are multiple shapes connected to the same shape as its sibling, at one or both ends. They are laid out all on the same side as their cluster siblings.

For example, notice that "a" and "c" could have minimized distance by going to "x"'s left and right sides, but chose to remain clustered together instead.




Figure 2: Example of a cluster

5.3 Hierarchy

Hierarchical structures are identified by multiple levels of shapes connected in the same direction and undergoes layout with a special set of rules (e.g. even spacing between levels). Uniquely, TALA handles containers within hierarchies.




Figure 3: Example of a hierarchy with containers

5.4 Balanced connections

When routing connections, TALA tries to target ports that make them look balanced, given other connections. For example, if it's the only connection on that side, it will try to route to the center – if 3, then at $1/3$ intervals, etc.




Figure 4: Example of connection balancing

5.5 Dynamic label positioning

Label positions are fixed everywhere. They are positioned to where they avoid obstructions like other shapes and connections.

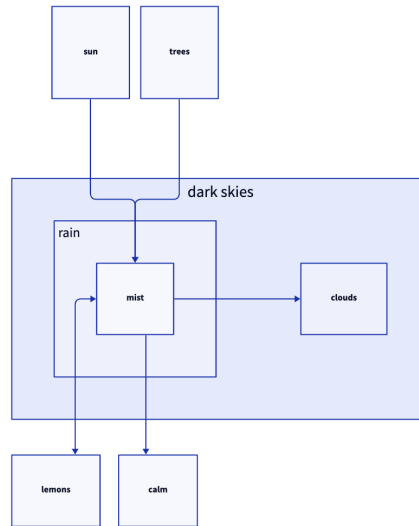


Figure 5: Shape label going to top-left to not collide with routes

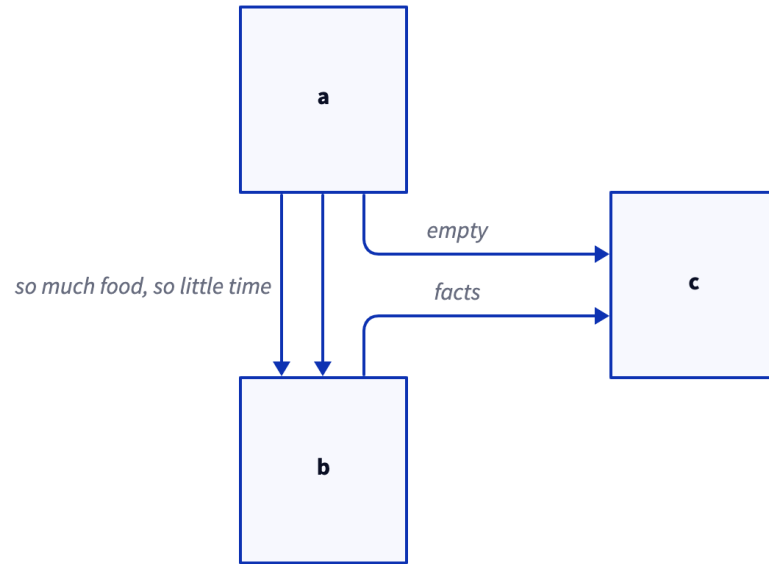


Figure 6: Connection labels avoiding each other

5.6 Square aspect ratio

Non-connected subgraphs are bin-packed to prefer a square aspect ratio.

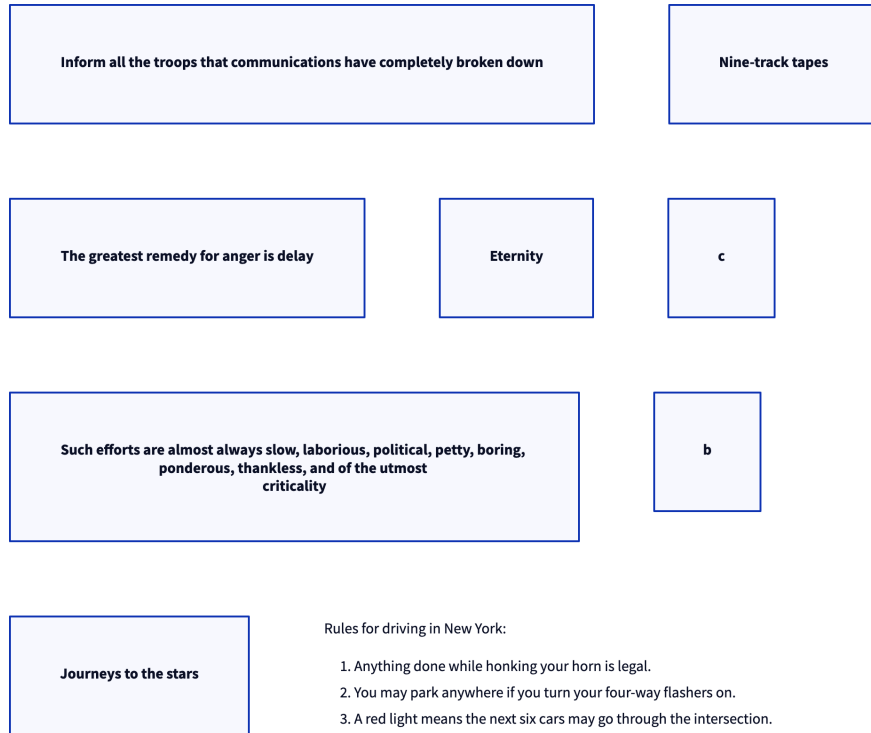


Figure 7: Example of bin-packing

5.7 Direction

Unlike other layout directions, TALA is able to specify **direction** keyword per container level. The default direction puts a slight preference on down-right (connections point down or right), though if you specify, the preference is stronger.

This keyword is a suggestion, not a requirement. Some connections may still flow the other direction when it improves other heuristics. But in general, the children of container will prefer the direction specified.




Figure 8: Top level direction is down, container a's direction is left, container b's direction is up, container c's direction is right

5.8 Near

On top of constants (e.g. `top-center`), `near` can target other shapes in TALA. Provide the absolute ID (e.g. a shape `world` inside of container `hello` is specified as `hello.world`) to the `near` property to tell TALA to place it in proximity of that shape.




Figure 9: Example of the near attribute

5.9 SQL table column matching

In TALA, connections between two `sql_table` shapes that specify the column will route to that specific column. These diagrams are also known as Entity-Relation Diagrams (ERDs).

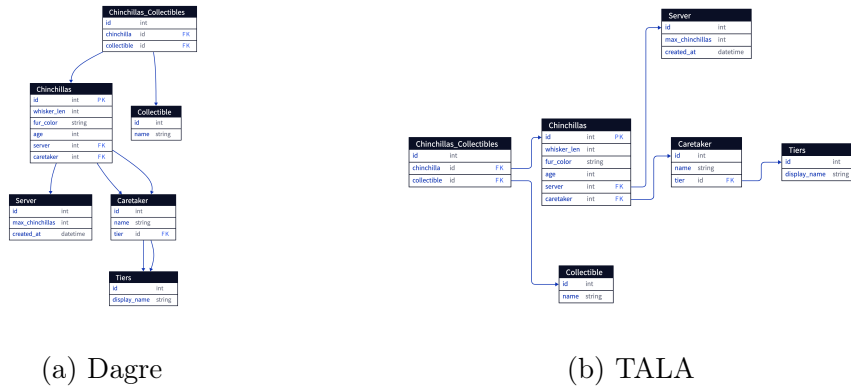


Figure 10: Comparison of ERD routing in Dagre vs TALA

5.10 Sequences

If you connect a chain of shapes with their type set to **step**, TALA arranges them tail to head instead of using connections as other shape types would.




Figure 11: Example of multiple step shapes to form sequence

5.11 Self connections

A self-connection is $x \rightarrow x$. In TALA, one corner is reserved for one type of connection (there are 3 types: directed, undirected, bidirected). If there are multiple connections of the same type, the corner is reused.

Notice how the outer loop makes room for the inner loop's label.

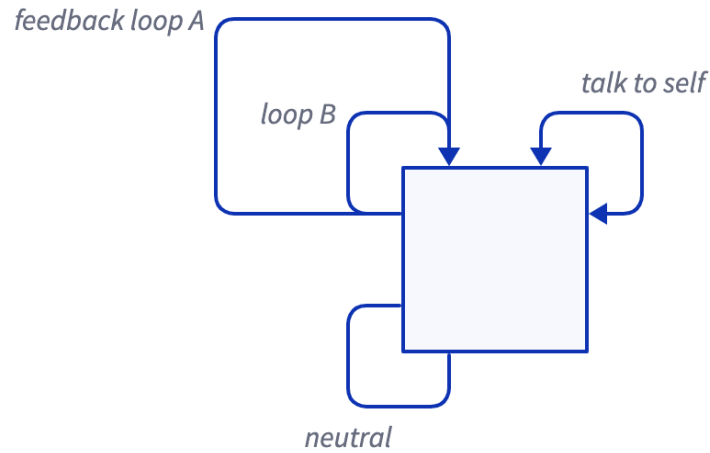


Figure 12: Example of multiple self-connections of varying relationships

6 Seed

Optimal placements of nodes that minimizes distance and crossings is NP-hard. More so when constraints like hierarchical are removed. TALA searches with heuristics to approximate. This search space has some randomness at each step. Choosing a different seed for this randomness can have significant impact on the overall layout, as it may converge on an entirely different look at the end.

The Terrastruct web app takes 3 seeds and runs them all in parallel, and uses a different set of heuristics to choose the best one. On the CLI, you can specify the seed with a flag, `tala-seed`.

```
d2 --layout tala --tala-seed 44 input.d2
```

Listing 2: Specifying a seed on the CLI

The intended usage is to "re-roll" for a better layout if TALA did not give an ideal layout with the default seed.

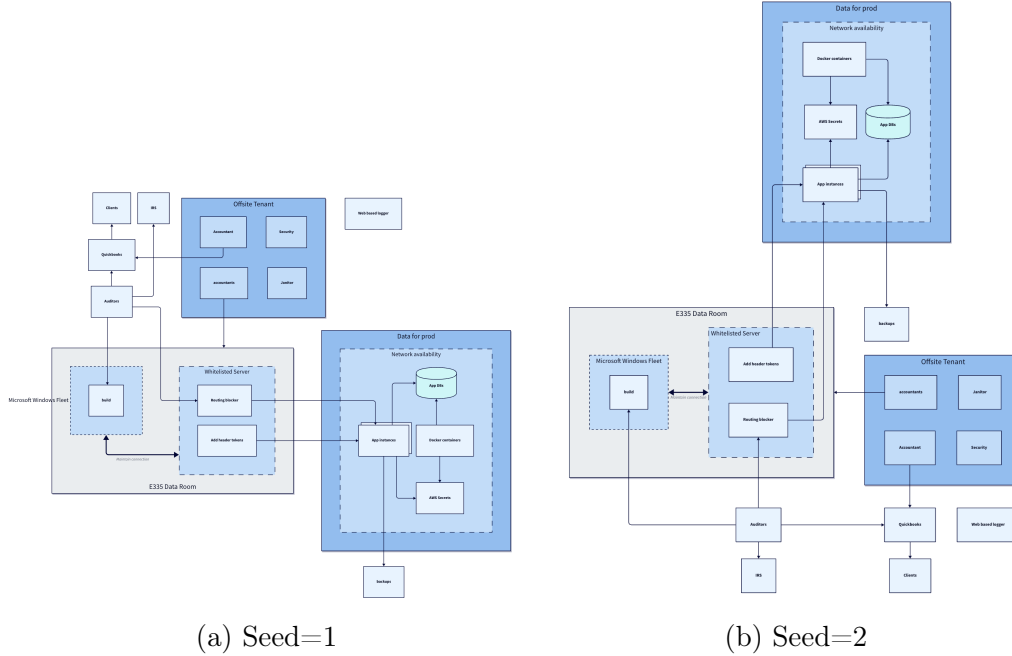


Figure 13: Same diagram, different seeds used

Not all diagrams will look different with different seeds – multiple random seeds may converge to the same diagram.

7 Tips

7.1 Dimensions

By default, the dimension of a shape is the dimension of its label, with some padding. TALA does not currently change dimensions other than to expand containers enough to fit children and make cluster siblings uniform. It won't know if you intended a shape to have a larger presence than others. For example, a central "hub" that many connect to. To achieve layouts that more closely match the ideal, you should supply **width** and **height** to shapes where you imagine it to be larger or smaller in scale than others.

Also, connections may look worse bending and twisting to try to reach a well-connected shape situated in the middle of obstructions. Giving custom

dimensions can give more space for routes to connect to.

7.2 Label collisions

Though labels dynamically avoid collisions, it will not influence the layout to make room for a disproportionately sized label. For a long label, you may find better results splitting them up into multiple lines with `\n`.

8 FAQ

8.1 Will TALA require an internet connection?

No, TALA does not collect telemetry or use the internet in any way except to ping to check the status of a license. This is only done when necessary, e.g. if you purchased a month subscription, TALA will ping at the start of the next month and renew automatically if the subscription is ongoing. If you purchased a year, it won't ping for a year. The only data that's sent in these pings is the API token itself. No diagrams or anything else leaves your computer.

8.2 Can I use TALA for diagrams which are not software architecture?

Yes, and it's known to give good results in many cases. But there's two things to note.

1. There are many categories of diagram types which are "solved". For example, for mind maps, the algorithm that mind-map-specific software has no improvements left to be made. If your diagram type falls under one of these categories, you'll definitely get better results with specific software. Perhaps D2 will support these, but it will be outside of TALA (e.g. D2's sequence diagram layout engine is bundled in D2).
2. TALA is tested exclusively on software architecture diagrams. It is the focus, and we will not change the algorithm to better handle any other type.

8.3 Can I use TALA to layout separately from D2?

Yes, it's possible – they communicate with a serialized, JSON representation of a diagram. There's no reason you couldn't build that JSON from somewhere

other than D2. However, it's not an intended use case, so isn't well-documented. If you have a use case in mind, please let us know and we'll help you.