

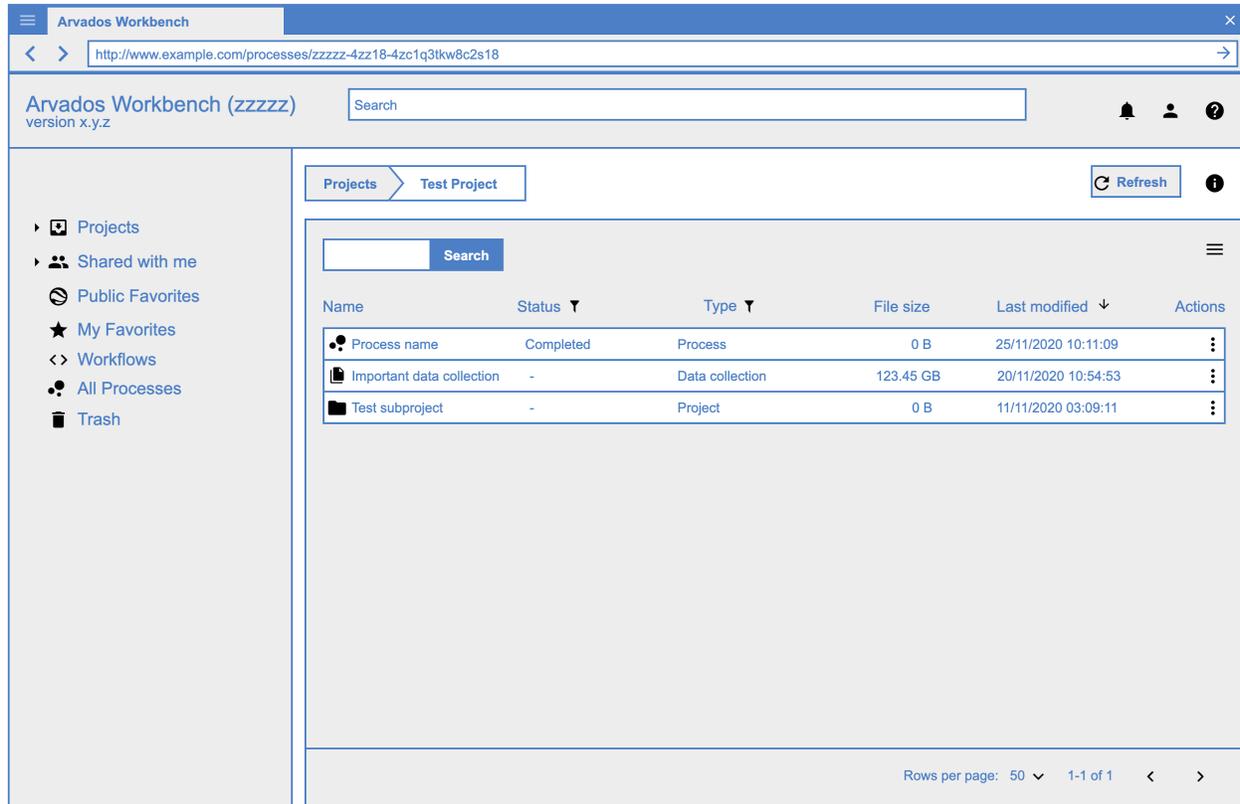
# General design guidelines

Due to the great amount of data needed to be displayed, and the assumption that showing it all at once is not a very good way of doing easy-to-use interfaces, I've tried to group the data following 3 different user *personas*:

1. **Monitor:** A user that is checking the running workflow.
  - a. Views live logs.
  - b. Checks workflow progress.
  - c. Checks if some step failed or some warning is present.
  - d. Cancels and re-runs processes.
2. **Debugger:** A user that's trying to find the reason for a failed workflow.
  - a. Checks inputs (command or workflow) and outputs.
  - b. Checks resulting logs.
  - c. Checks container image used.
  - d. Checks runtime constraints & scheduling parameters.
  - e. Searches for failed steps.
3. **Viewer:** A user that wants to check a successful workflow and its results.
  - a. Checks run time and time scaling.
  - b. Shares workflow output.
  - c. Checks which workflow the container was launched from.

# Project view (listing)

## Current version

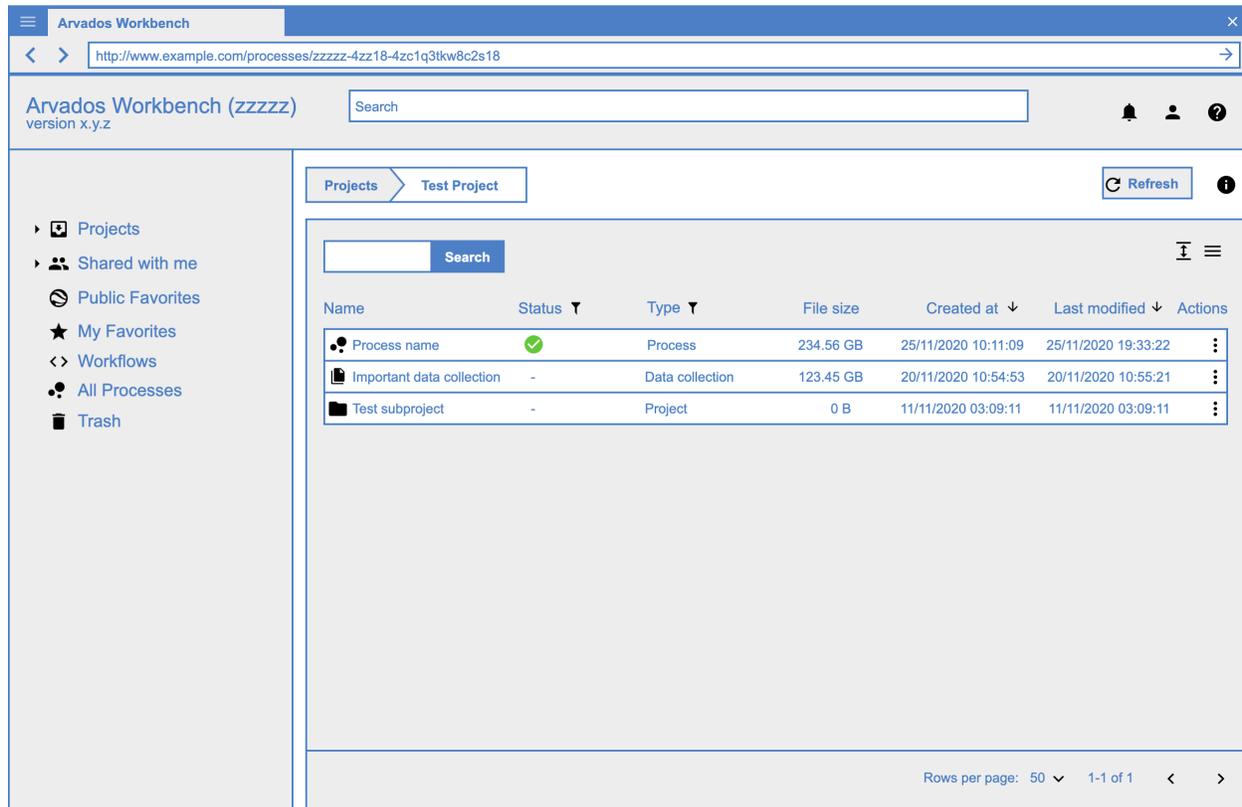


The current UI design displays the project as just a folder with objects inside it. Properties and description are second class citizens limited to its details panel. There's room for improvement there but in this case we'll focus on the Process UI.

The processes are listed as items with limited amount of information, requiring the use to open the details panel (the "i" icon) to get more data about it.

# Proposed version

## Listing



The proposed project listing includes:

- Process' status indicators as icons to save on horizontal space
- An "expand" button at the top-right corner to make the listed rows show additional information. We'll start by making processes rows expandable and can iterate further on other types in the future.
- Addition of "created at" sortable column. If horizontal space is not enough, we could make some columns hidden by default.
- "File size" column showing the process' output collection size. Also, we could rename the column to just "Size".

## Listing with extended details panel

The screenshot shows the Arvados Workbench interface. The top navigation bar includes the Arvados Workbench logo and a search bar. The main content area is divided into a sidebar on the left with navigation options like 'Projects', 'Shared with me', and 'Public Favorites'. The central area displays a table of processes under the 'Test Project' tab. The table has columns for Name, Status, Type, and File size. The first row, 'Process name', is highlighted in light blue and has a green checkmark in the Status column. To the right of the table, a 'Process name' details panel is expanded, showing various metadata fields and sections for 'Runtime constraints' and 'Scheduling'.

Name	Status	Type	File size
Process name	✓	Process	234.56 GB
Important data collection	-	Data collection	123.45 GB
Test subproject	-	Project	0 B

**Process name** DETAILS

Process UUID  
zzzzz-xvhdp-4zc1q3tkw8c2s18

Owner  
Test Project (zzzzz-j7d0g-a8iau300x48obzj)

Requesting Container UUID  
(none)

Container UUID  
zzzzz-dz642-231hdt9w8c54w2q

Started at  
MM/DD/YYYY HH:MM:SS

Run time  
XXhYYmZZs (AAhBBmCCs queued)

Logs  
log-collection-uuid

Outputs  
output-collection-uuid

Container image  
99b0201f4cade456b4c9d343769a3b70+261

Command  
[echo, 'hello', 'world']

**Runtime constraints**

```
API true
keep_cache_ram 268435456
ram 1073741824
vcpus 1
```

**Scheduling**

```
preemptible false
max_run_time 0
```

Whenever the user clicks on a row, the details panel expands to show the most important information about the process. Right now the user has to click on the “i” icon to get the panel extended.

JSON information like Runtime constraints and Scheduling parameters should be formatted appropriately for ease of reading.

UUIDs should be clickable and have a “copy to clipboard” button.

## Listing with expanded rows

The screenshot shows the Arvados Workbench interface. The top navigation bar includes the Arvados Workbench logo and a search bar. The main content area is divided into a left sidebar with navigation options (Projects, Shared with me, Public Favorites, My Favorites, Workflows, All Processes, Trash) and a main panel. The main panel shows a breadcrumb trail for 'Test Project' and a 'Refresh' button. Below this is a table listing processes. The first row is expanded, showing detailed information about a process, including its UUID, owner, container details, and runtime constraints.

Name	Status	Type	File size	Last modified	Actions
Process name	✓	Process	234.5.6 GB	25/11/2020 10:11:09	⋮
<div>Process UUID zzzzz-xvhdp-4zc1q3tkw8c2s18</div> <div>Owner Test Project (zzzzz-j7d0g-a8iau300x48obzj)</div> <div>Requesting Container UUID (none)</div> <div>Container UUID zzzzz-dz642-231hdt9w8c54w2q</div> <div>Started at MM/DD/YYYY HH:MM:SS</div> <div>Run time XXhYYmZZs (AAhBBmCCs queued)</div> <div>Logs log-collection-uuid</div> <div>Outputs output-collection-uuid</div> <div>Container image 99b0201f4cade456b4c9d343769a3b70+261</div> <div>Command [echo, 'hello', 'world']</div> <div>Runtime constraints API true keep_cache_ram 268435456 ram 1073741824 vcpus 1</div> <div>Scheduling preemptible false max_run_time 0</div>					
Important data collection	-	Data collection	123.45 GB	20/11/2020 10:54:53	⋮
Test subproject	-	Project	0 GB	11/11/2020 03:09:11	⋮

When the user clicks on the “expand rows” button, the details panel collapses itself (if needed) and all expandable rows show additional information like the mockup above.

The idea is to be able to use the browser’s “Find” feature to quickly scan the entire listing.

Part of this expanded information UI can be re-used from the “details panel”.

If the user requires to only view expanded processes from within a project, they can filter them by using the already existing Type column filtering.

# Process view

## Current version

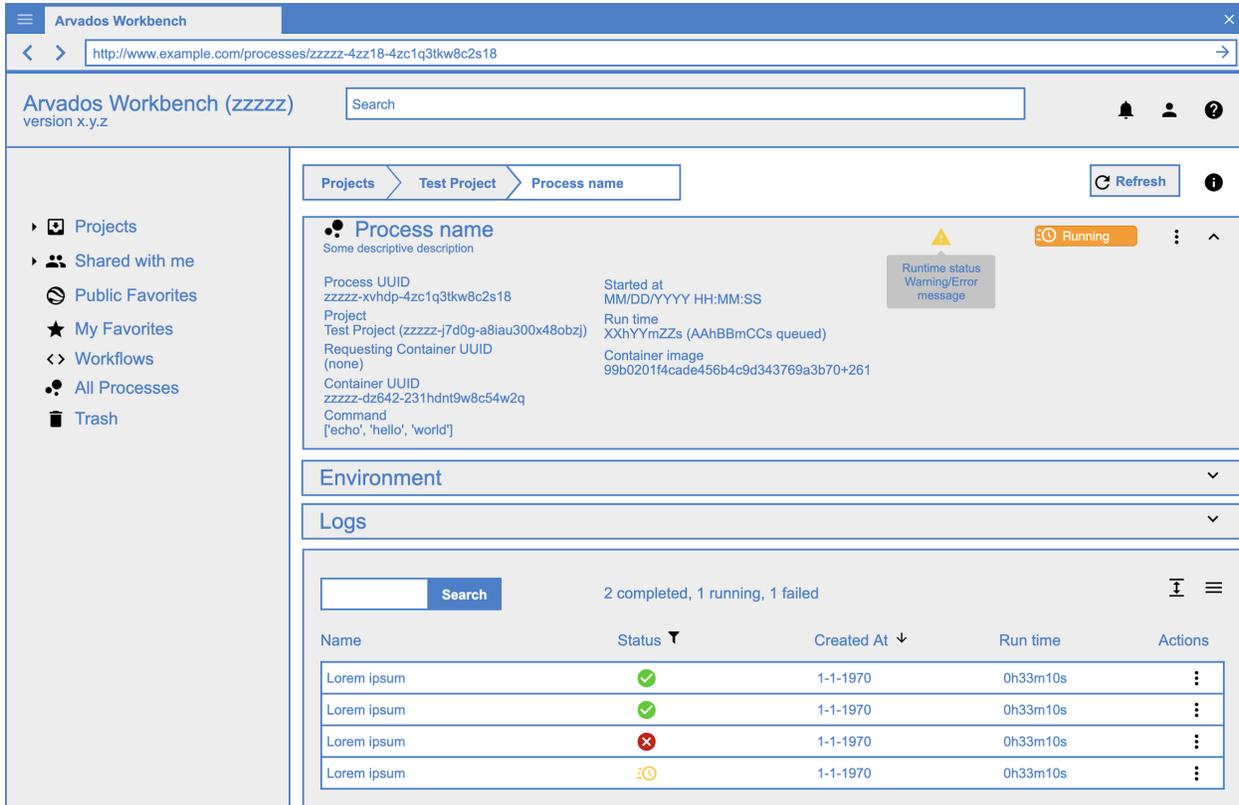
The screenshot shows the Arvados Workbench interface. The top navigation bar includes the text "Arvados Workbench" and a search input field. Below this, a breadcrumb trail shows "Projects" > "Test Project" > "Process name". A "Refresh" button is visible next to the breadcrumb. The main content area is divided into two sections. The top section displays the process name "Process name" with a "Completed" status indicator. Below this, there is a diagram showing "From" and "To" connections, with "Outputs" and "Inputs" labels. The bottom section features a search bar and a table of process instances.

Name	Status	Created At	Run Time	Actions
Lorem ipsum	Completed	1-1-1970	0h33m10s	⋮
Lorem ipsum	Completed	1-1-1970	0h33m10s	⋮
Lorem ipsum	Completed	1-1-1970	0h33m10s	⋮
Lorem ipsum	Completed	1-1-1970	0h33m10s	⋮

The current process view doesn't give the user useful information up front, but hides too much of it behind the "details panel" sometimes displaying data points incorrectly or in not so efficient ways.

# Proposed version

## General info panel



Due to the amount of information available to display and thinking about grouping them following the design guidelines described above, the proposed UI includes 3 sections: the general information panel, the environment panel and the live logs viewer.

These sections could be expandable/collapsible individually or only just one being able to expand at a time.

To save vertical space we could also divide these sections via a tab layout like workbench1 does.

The general information panel lists the data most likely to be needed by the “Viewer” and “Monitor” *personas*, being able to quickly see if there’s some warning or error via a runtime status icon next to the process’ state indicator.

The steps listing could also maintain the same behavior as the project listing: providing an “expand rows” button to show additional information on every step. Probably this would be trivial to implement by reusing the UI components.

# Live logs panel

The screenshot displays the Arvados Workbench interface. At the top, the browser address bar shows a URL: `http://www.example.com/processes/zzzzz-4zz18-4zc1q3tkw8c2s18`. The main header includes the text "Arvados Workbench (zzzzz)" and "version x.y.z". A search bar is present next to it. On the left, a navigation sidebar lists: "Projects", "Shared with me", "Public Favorites", "My Favorites", "Workflows", "All Processes", and "Trash". The main content area shows a breadcrumb trail: "Projects > Test Project > Process name". A "Refresh" button is located to the right. Below the breadcrumb, the process name "Process name" is displayed with a yellow warning icon and a "Running" status indicator. The "Environment" section is collapsed. The "Logs" section is expanded, showing a terminal window with the following log output:

```
2021-03-16T14:45:23.438992044Z crunch-run 2.2.0-dev20210312200937 (go1.13.4) started
2021-03-16T14:45:23.445145544Z Executing container 'ce815-d2642=mlvthakallezpx41'
2021-03-16T14:45:23.445170245Z Executing on host 'compute-99cb150b26149780de44b929577e1aed-e4cysk8im73bdw3'
2021-03-16T14:45:23.654170059Z Fetching Docker image from collection '960dcf39a7a7cc3e345af955a4308930+261'
2021-03-16T14:45:23.680590721Z Using Docker image id 'sha256:b228ffe7458b89eb1bc095a118a053750f64f892a1fcedd5c090f91e6f544b2a'
2021-03-16T14:45:23.683484915Z Loading Docker image from keep
2021-03-16T14:45:23.654419568Z notice: reading stats from /sys/fs/cgroup/cpuacct/cgroup.procs
2021-03-16T14:45:23.654462269Z notice: reading stats from /sys/fs/cgroup/memory/memory.stat
2021-03-16T14:45:23.654664776Z mem 527515648 cache 0 swap 1386 pgmajfault 195604480 xss
2021-03-16T14:45:23.654687276Z notice: reading stats from /sys/fs/cgroup/cpuacct/cpuacct.stat
2021-03-16T14:45:23.654745178Z notice: reading stats from /sys/fs/cgroup/cpuset/cpuset.cpus
```

Below the logs, a summary bar indicates "2 completed, 1 running, 1 failed". A search bar is provided. A table lists the process history:

Name	Status	Created At	Run time	Actions
Lorem ipsum	✓	1-1-1970	0h33m10s	⋮
Lorem ipsum	✓	1-1-1970	0h33m10s	⋮
Lorem ipsum	✗	1-1-1970	0h33m10s	⋮
Lorem ipsum	🔄	1-1-1970	0h33m10s	⋮

Currently, live logs are accessible via the “actions menu” (3-dotted icon), changing the entire view for that purpose. I believe it can be a lot more comfortable for the user to not have to go back and forth changing panels when needing to quickly check how the process is running. Another solution would be to make the live log viewer appear on a quickly dismissable modal window.

## Environment panel

The screenshot shows the Arvados Workbench interface. The top navigation bar includes the Arvados Workbench logo and a search bar. The main content area is divided into a left sidebar with navigation options (Projects, Shared with me, Public Favorites, My Favorites, Workflows, All Processes, Trash) and a main panel. The main panel shows a breadcrumb trail: Projects > Test Project > Process name. Below this, there's a 'Process name' header with a status indicator 'Running' and a 'Refresh' button. The 'Environment' section is expanded, showing three sub-panels: 'Inputs' (listing /path/to/input-file-1 and /path/to/input-file-2), 'Outputs' (listing zzzzz-4zz18-4zc1q3tkw8c2s18), 'Command' (a JSON array of container options), 'Runtime constraints' (a JSON object with API, keep\_cache\_ram, ram, and vcpus), and 'Scheduling' (a JSON object with preemptible and max\_run\_time). Below the environment panels is a 'Logs' section with a search bar and a table of log entries. The table has columns for Name, Status, Created At, Run time, and Actions. It shows four entries, each with a status icon (green checkmark, red X, or yellow warning) and a run time of 0h33m10s.

The *Debugger persona* may need to have all the environment information in one place, so this panel would show inputs (from command and/or workflow), outputs, runtime constraints and scheduling parameters.

The proposed layout tries to fit often size changing data in the sense horizontal and/or vertical space. Vertical scrollbars can be added when the listings are too long.

Inputs and outputs should be clickable links and also any keep reference being displayed on the Command sub-panel, as we do in workbench1.

## Additional ideas

- Any collapsible panel could have a bit of the more crucial information, for example:
  - The General info panel could have in addition to the process' name and status, the run time.
  - The Logs panel when collapsed could include a clickable UUID link to the process' log collection.
  - The Environment panel when collapsed could show the container's command in a oneliner format, allowing us to remove it from the General info panel.
- If the segmented process view is not good enough, we could default to workbench1's way of showing the data.

- Right now the steps listing's bottom horizontal scroll bar goes past the browser's viewport, we could correct that to behave like the project listing.