

# Enabling Scientific Communities through ParaView

Code as Community

J. Fausty, M. Westphal and F. Mazen, PASC23

# Kitware / Leader in AI & scientific open source solutions

## Software development

Based on open source tools  
300+ active projects worldwide



## Sustained Growth

Since creation of the company  
100% employee-owned



## 230 employees Worldwide

6 offices across USA/Europe



## 65% staff with PhD or Master

High Level customer expertise



## 20+ years of expertise

Kitware USA, 1998  
Kitware Europe, 2010



## Revenue 2020

\$39M consolidated



# Customers / Various fields of application

## Academics

70+ academic institutions worldwide

## Government agencies

50+ government agencies and national laboratories

## Commercial companies

Over 500 commercial customers

## Medical

Image processing, multimodal visualization, image registration & segmentation, assisted surgery, custom software...

## Energy

HPC, in-situ simulation, scientific visualisation, particle flow, fluid mechanics, ground exploration...

## Intelligence

Scene analysis, big data analysis, scientific visualization, flow analysis...



# Kitware / **Services**



**TRAINING**



**SUPPORT**



**DEVELOPMENT**



**GRANT  
COLLABORATION**

# Open Source Benefits / Shifting Power



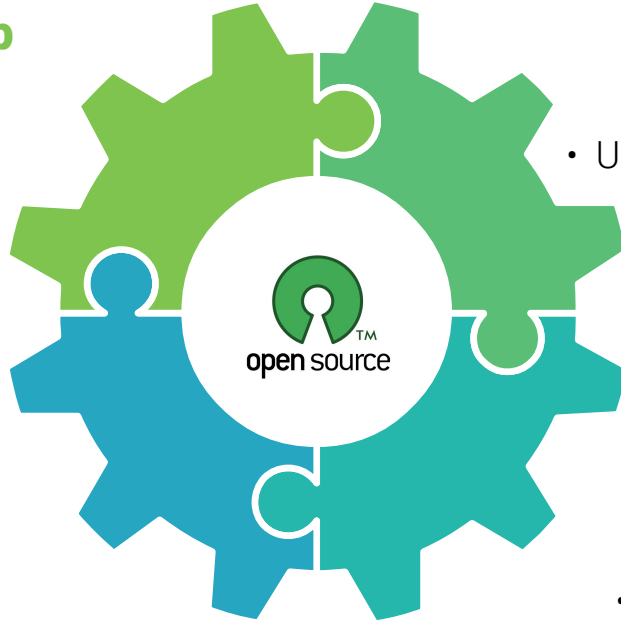
## Source code ownership

- Source code ownership
- Integration with commercial software solutions



## Cost effectiveness

- No license fee
- No vendor lock-in
- Shared maintenance costs



## Flexibility and Agility

- Continuous development
- Up to date with new technologies
- Ability to customize and fix



## Security

- Robust software and libraries
  - Transparency
  - Community effort
- Open Innovation mitigates risk

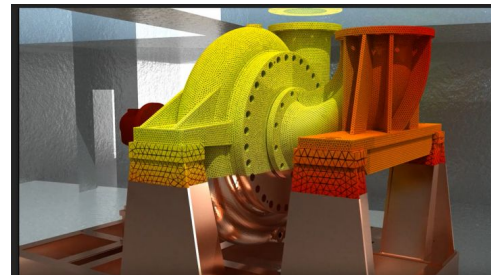
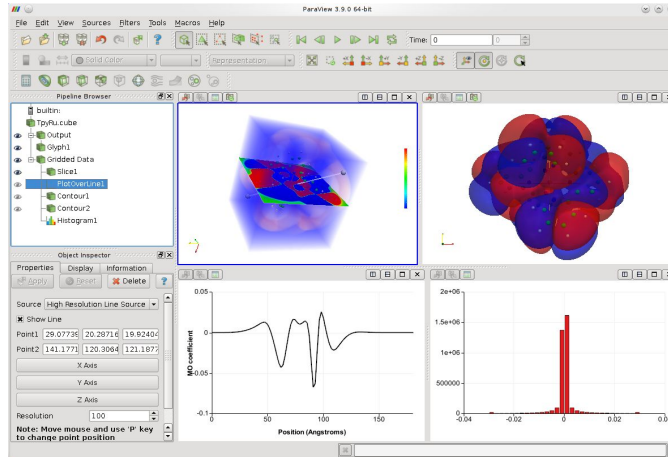
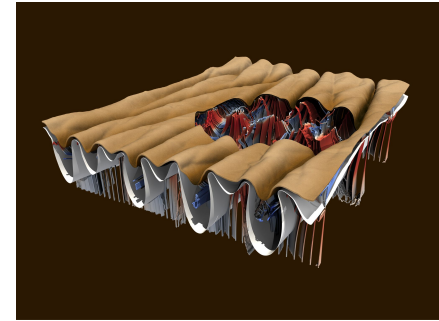


# What is ParaView?

An application and architecture for display and analysis of massive scientific datasets.

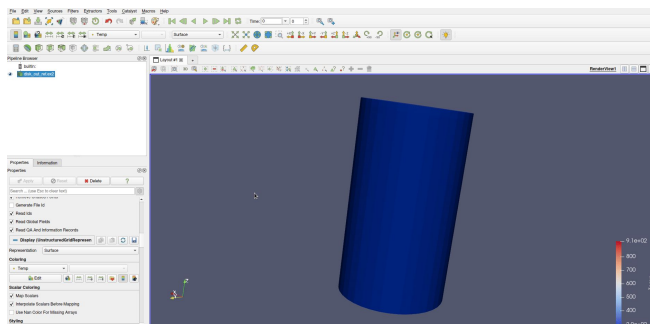
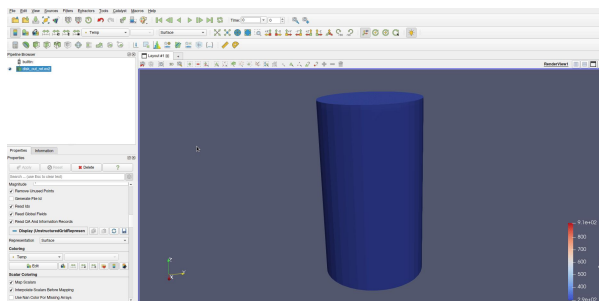


- **Open-source**, multi-platform, data analysis and **visualization** application
- Analysis of **extremely large datasets** using **distributed computing** resources
- **Advanced rendering capabilities** such as physically based rendering and ray tracing solutions



# Quick demos

## Clipping

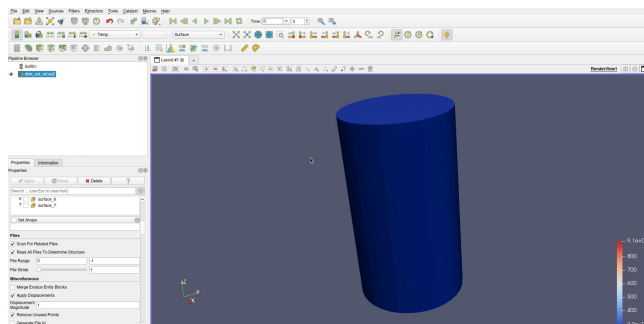


## Volume Rendering

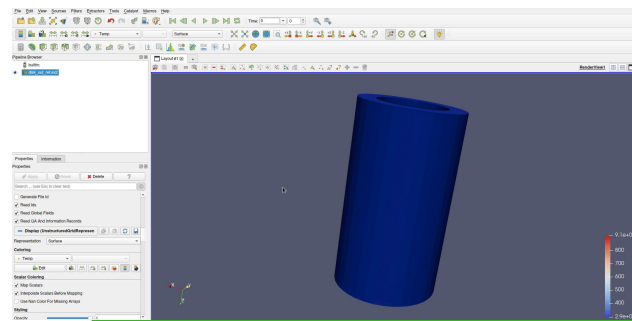


... and much more

## Streamlines



## Contouring



## Mesh Inspection

# Some open-source ParaView facts and resources

- **6 month release** schedule
- **160 190 commits** (~ 3500 commits yearly basis)
- **349 contributors** (~20/month)
- Code hosted on a **GitLab** instance
- **~400K downloads/year**
- First public release in **2002**

<b>Codebase (LOC)</b>	<b>1,663,581</b>
<b>Effort (est.)</b>	<b>469 Person Years</b>

Community	Members	Topics per month
VTK Discourse	3000	100
ParaView Discourse	4500	100

ParaView GitLab: <https://gitlab.kitware.com/paraview/paraview>  
Stats on the repository: <https://www.openhub.net/p/ParaView>  
Discourse forum link: <https://discourse.paraview.org/>

# Identification of ParaView's Community Roles

User and Contributor Profiling

# Who constitutes ParaView's community?

## Users

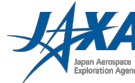
### Academia



... and many more



### Industry



... and many more

## Contributors

- **Kitware** (~98% of dev activity)
- **Professional Contributors\*** (~ 1-2% of dev activity)
- **Hobbyists** (< 1% of dev activity)
- ... but there is more to contributing than development

\*contributors who do so as part of their work

# Some Community User Profiles



**Maria, Ph.D.**

**Occupation**

Postdoctoral Researcher

**ParaView Usage**

Visualization of the results of her CFD simulation data on wind turbine optimization.

Desktop  
End User

Scripter



**Hugo, Engineer**

**Occupation**

Simulation and IT Support

**ParaView Usage**

Provide and develop curated ParaView deployments and workflows internally at XXX.

Deployer

Plugin  
Maintainer



**Sarah, Engineer**

**Occupation**

GPU Driver Development

**ParaView Usage**

Stress test, benchmark and evaluate new HPC GPU driver implementations using ParaView.

API  
Consumer

Benchmark

# Some Community Contributor Profiles



**Takashi, M.Sc.**

**Occupation**

Ph.D. Student

**ParaView Usage**

Volume rendering of astronomy simulations of galaxy dynamics.

Community  
Support

Issue  
Opener



**Anabelle, Engineer**

**Occupation**

Automotive Acoustics Engineer

**ParaView Usage**

Visualization and automatic post-treatment of acoustic simulations of car geometries.

Issue  
Opener

Funder



**Mathew, Engineer**

**Occupation**

Senior Software Engineer at  
Kitware

**ParaView Usage**

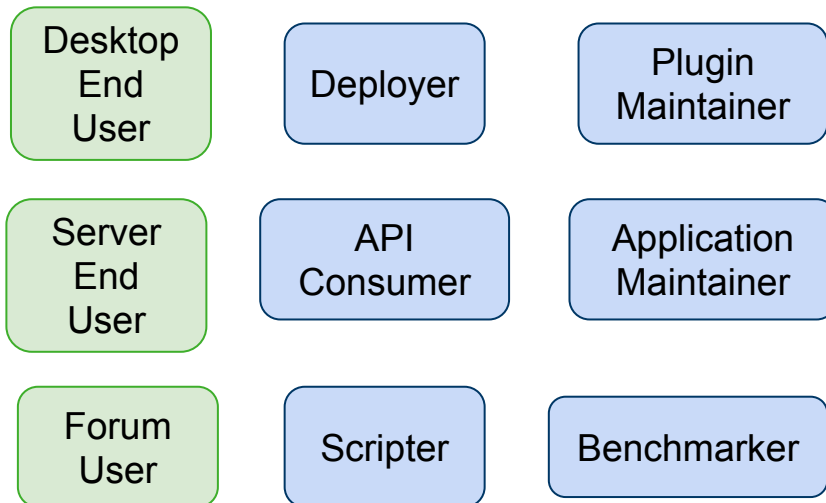
Develops and maintains  
ParaView on a daily basis.

Developer

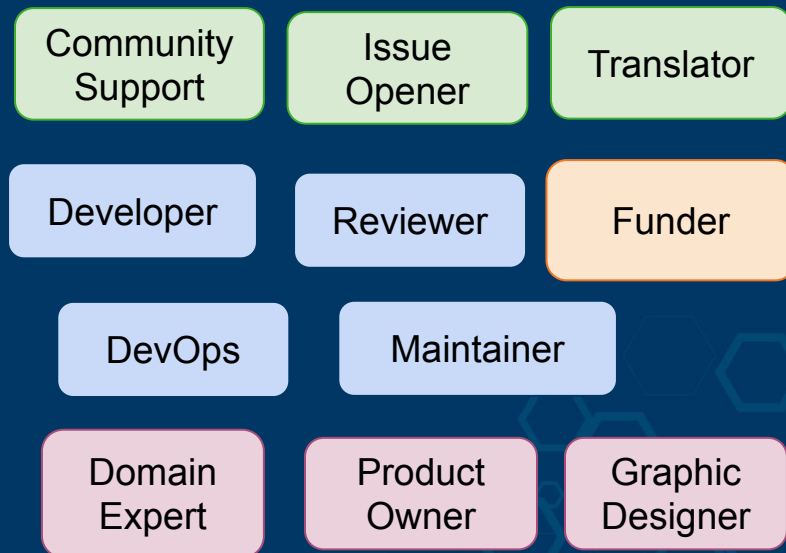
Maintainer

# A bird's eye view of community roles

## Users



## Contributors



# ParaView's development drivers and feedback loops

Transient VTKHDF - Case Study

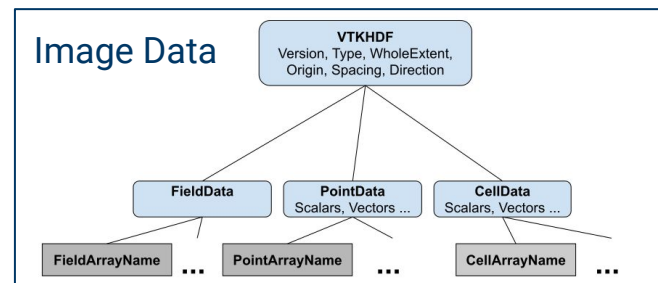
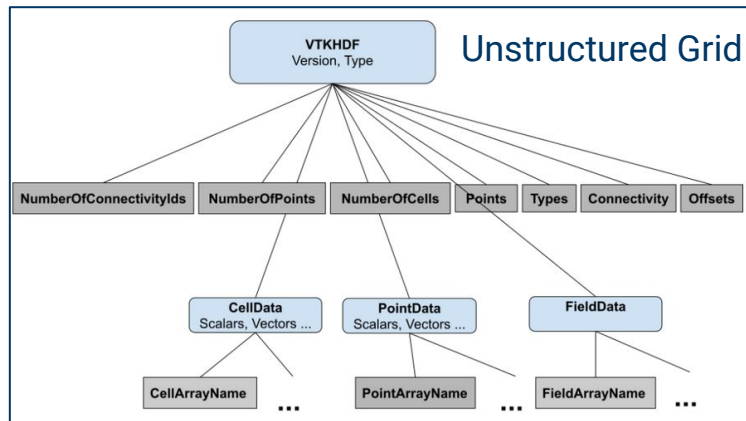
# ParaView's development drivers and feedback loops: Case Study

Funder

## Funder Goals:

- Extend the existing VTKHDF file format specification to handle transient data
- Add the transient data support to the associated reader in VTK/ParaView

Existing designs:



# ParaView's development drivers and feedback loops: Process

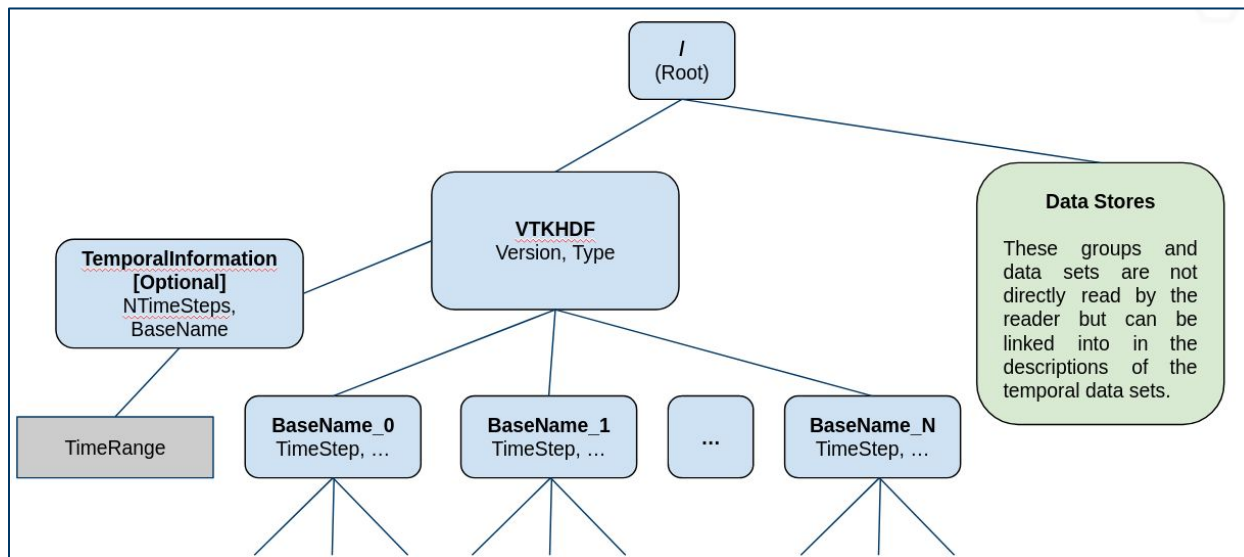
wants to  
develop a  
feature

Funder

# ParaView's development drivers and feedback loops: Case Study

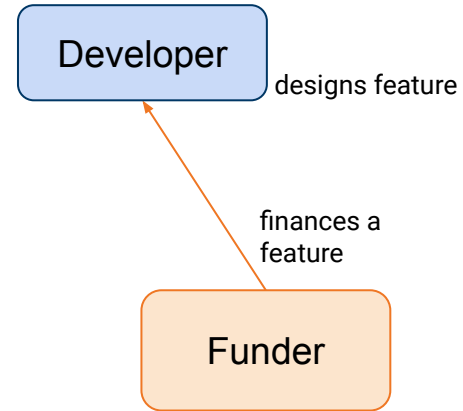
## Initial Design:

Developer



# ParaView's development drivers and feedback loops:

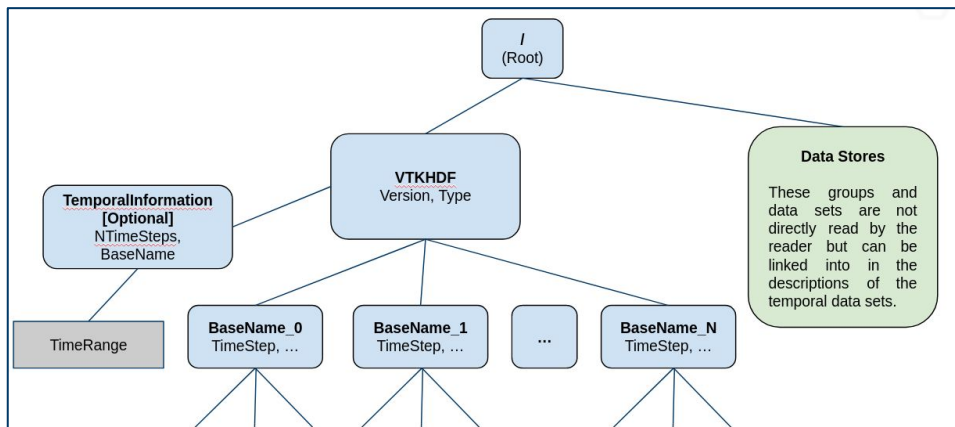
## Process



# ParaView's development drivers and feedback loops: Case Study

## Initial Design:

Developer



## Design Review Conclusions:

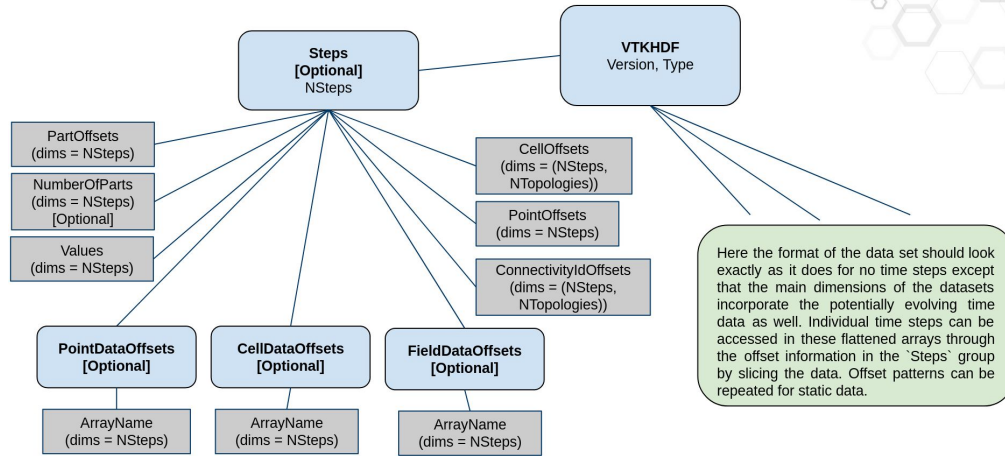
- Too much meta-data in file which might hurt parallel IO performance

Domain Expert

Product Owner

# ParaView's development drivers and feedback loops: Case Study

## Iterated Design



Developer



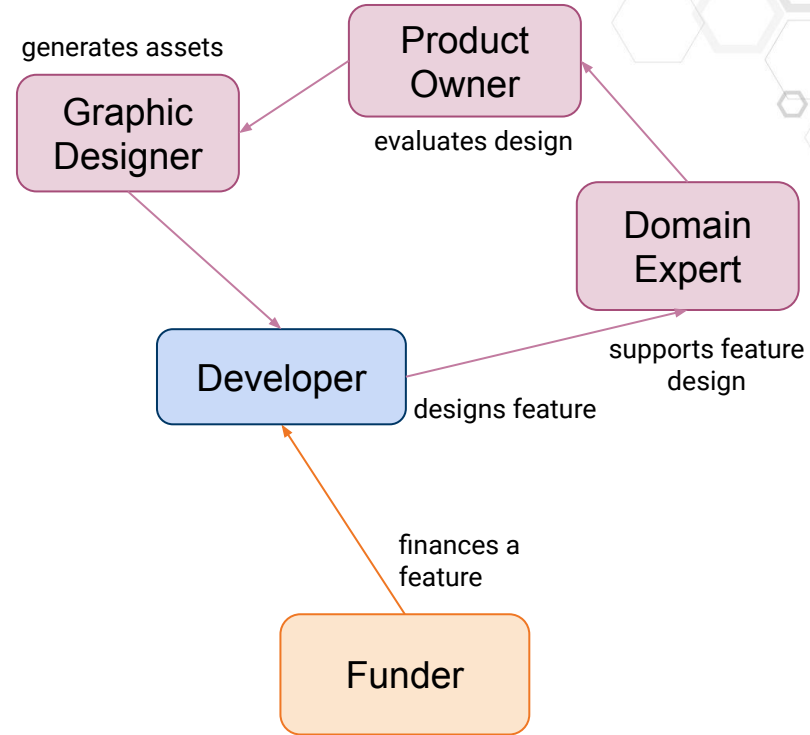
Domain Expert



Product Owner

# ParaView's development drivers and feedback loops:

## Process

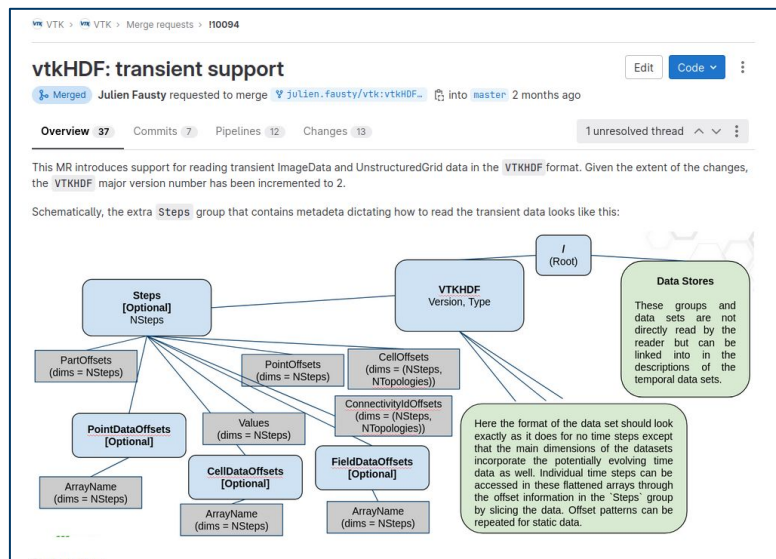


# ParaView's development drivers and feedback loops: Case Study

## Implementation and Merge Request:

Reviewer

Developer



Overview 37 Commits 7 Pipelines 12 Changes 13 1 unresolved thread

Mathieu Westphal (Kitware) @mwestphal started a thread on an old version of the diff 2 months ago  
Resolved 2 months ago by Julien Fausty

IO/HDF/vtkHDFReader.cxx

```
807 808 {
809 809 }
810 810 return 0;
811 811 }
812 812 if (this->HasTransientData())
813 813 {
814 814     double* values = outInfo->Get(vtkStreamingDemandDrivenPipeline::TIME_STEPS());
815 815     if (outInfo->Has(vtkStreamingDemandDrivenPipeline::UPDATE_TIME_STEP()))
816 816     {
817 817         double requestedValue = outInfo->Get(vtkStreamingDemandDrivenPipeline::UPDATE_TIME_STEP());
818 818         this->Step = std::distance(
819 819             values, std::upper_bound(values, values + this->NumberOfSteps, requestedValue));
820 820         this->Step = this->Step - 1;
821 821     }
822 822 }
```

Mathieu Westphal (Kitware) @mwestphal · 2 months ago

this->Step >= this->NumberOfSteps ?

Julien Fausty changed this line in version 7 of the diff 2 months ago · [Compare changes](#)

Reply...

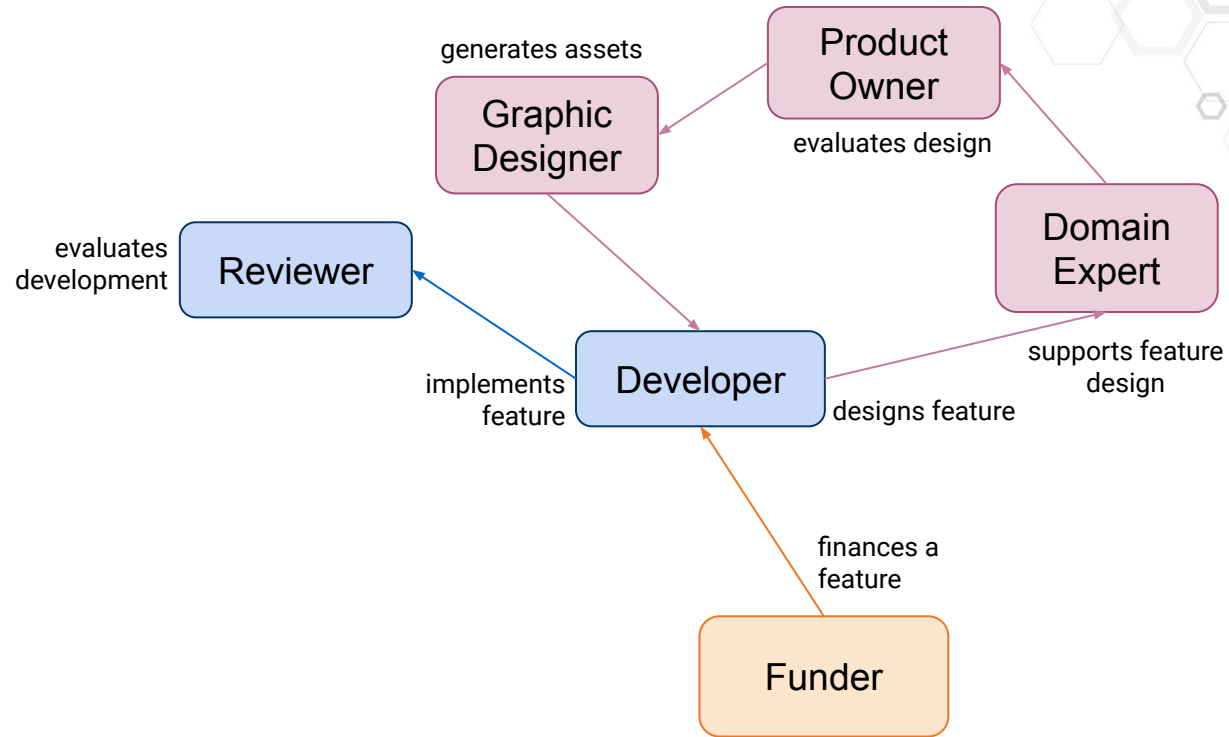
Unresolved thread

Mathieu Westphal (Kitware) @mwestphal · 2 months ago

+2 pending small comment and CI

# ParaView's development drivers and feedback loops:

## Process



# ParaView's development drivers and feedback loops:

## Case Study

## Merging and Continuous Integration

DevOps

vtk VTK > Merge requests > !10094

### vtkHDF: transient support

Merged Julien Fausty requested to merge [Julien Fausty/vtk:vtkHDF-](#) into [master](#) 2 months ago

Overview 37 Commits 7 Pipelines 12 Changes 13 1 unresolved thread

Status	Pipeline	Triggerer	Stages	Run pipeline
02:29:41	<a href="#">vtkHDFReader: fixing build warnings</a> #328345 18094 -> 8cc18ae6 <a href="#">latest</a> <a href="#">merge request</a>			
05:09:32 2 months ago	<a href="#">vtkHDFReader: fixing build warnings</a> #327989 18094 -> b2c87a46 <a href="#">merge request</a>			
05:41:07 2 months ago	<a href="#">vtkHDFReader: fixing build warnings</a> #327529 18094 -> bb7f8441 <a href="#">merge request</a>			
04:23:30 2 months ago	<a href="#">vtkHDFReader: fixing build warnings</a> #327508 18094 -> 84bcf80a <a href="#">merge request</a>			
In progress	<a href="#">vtkHDFReader: fixing build warnings</a> #327507 18094 -> d3e2b14a <a href="#">merge request</a>			

vtk VTK > Commits > 657c719a

Commit 657c719a authored 2 months ago by Julien Fausty Committed by Kitware Robot 2 months ago

Browse files Options

### Merge topic 'vtkHDF-transient-support'

[8cc18ae6](#) vtkHDFReader: fixing build warnings  
[539cfec2](#) vtkHDFReader: add documentation  
[882890ac](#) vtkHDF: add transient ImageData support  
[86fe342b](#) vtkHDF: add transient ImageData testing  
[fde03a1d](#) vtkHDFReader: place reverse extent ordering in image data section  
[8e2d6988](#) vtkHDFReader: add transient U6 support  
[de277bcb](#)

vtkHDF: add testing for transient U6 data

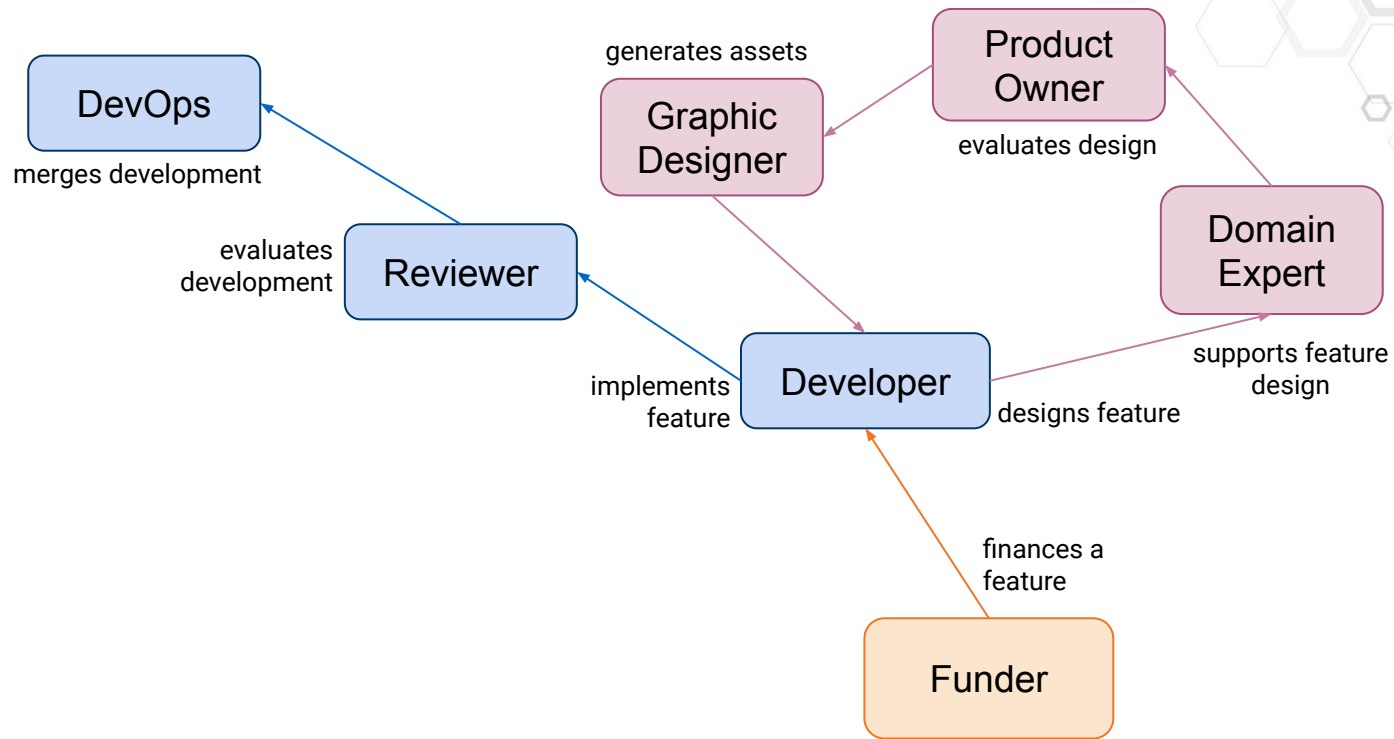
Acked-by: Kitware Robot <kwrobot@kitware.com>  
Acked-by: buildbot <buildbot@kitware.com>  
Reviewed-by: Thomas Galland <thomas.galland@kitware.com>  
Merge-request: !10094

parents [507f80ab](#) [8cc18ae6](#) [master](#)

1 merge request !10094 vtkHDF: transient support

Pipeline #328527 canceled with stages in 244 minutes and 25 seconds

# ParaView's development drivers and feedback loops: Process



# ParaView's development drivers and feedback loops:

## Case Study

## Deployment

Maintainer

DevOps

kitware.com contact us

about solutions resources companion tools customize Enter Keyword DOWNLOAD

download

### Download ParaView

**Get the Software**

You can either download binaries or source code archives for the latest stable or previous release or access the current development (aka nightly) distribution through Git. Specific license information can be found [here](#). This software may not be exported in violation of any U.S. export laws or regulations. For more information regarding Export Control matters please go to [https://kitware.com/export\\_control/index.html](https://kitware.com/export_control/index.html).

Version nightly

**ParaView**

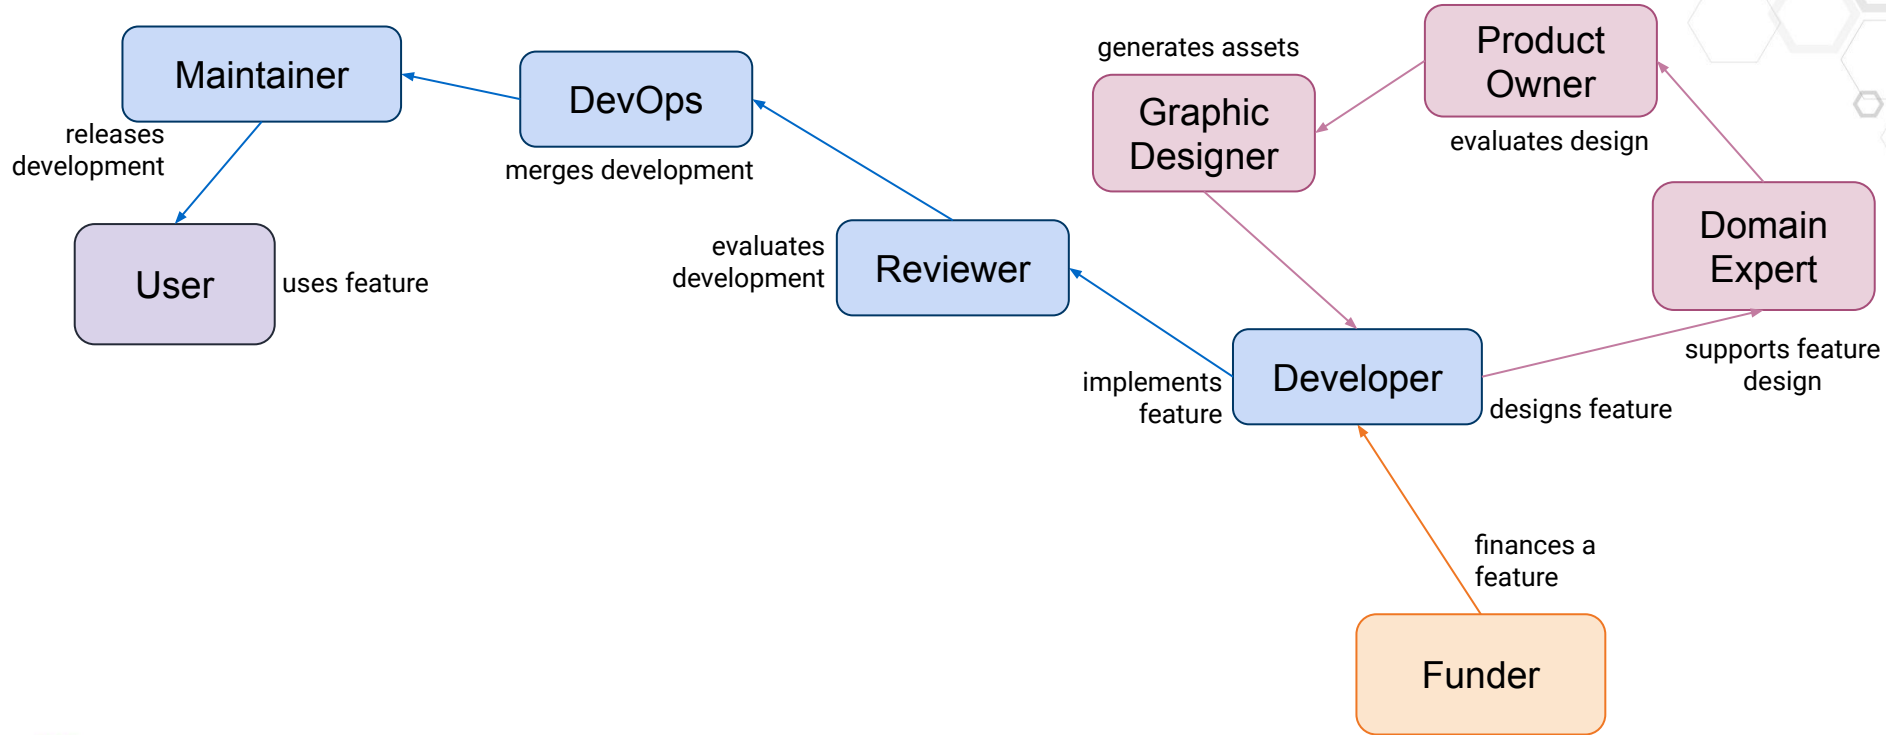
Full suite of ParaView tools, including the ParaView GUI client, pvpython, pvserver, pvbatch, and bundled MPI.

	Windows	Linux	macOS
ParaView-master-5.11.1-1284-g48ca581a4c-MPI-Linux-Python3.9-x86_64.tar.gz		2023-06-19 12:13	595M
ParaView-master-5.11.1-1282-ge8d1b6e973-MPI-Linux-Python3.9-x86_64.tar.gz		2023-06-19 04:41	595M
ParaView-master-5.11.1-1280-gd982f83420-MPI-Linux-Python3.9-x86_64.tar.gz		2023-06-19 01:37	595M
ParaView-master-5.11.1-1276-gef6e71a9a2-MPI-Linux-Python3.9-x86_64.tar.gz		2023-06-14 08:33	595M
ParaView-release-5.11.1-17-g7ef258d62c-MPI-Linux-Python3.9-x86_64.tar.gz		2023-06-14 07:00	557M
ParaView-master-5.11.1-1271-g522db06714-MPI-Linux-Python3.9-x86_64.tar.gz		2023-06-14 05:13	595M
ParaView-master-5.11.1-1262-g1d6b7dd8c1-MPI-Linux-Python3.9-x86_64.tar.gz		2023-06-14 01:30	595M

<https://www.paraview.org/download/>

# ParaView's development drivers and feedback loops:

## Process

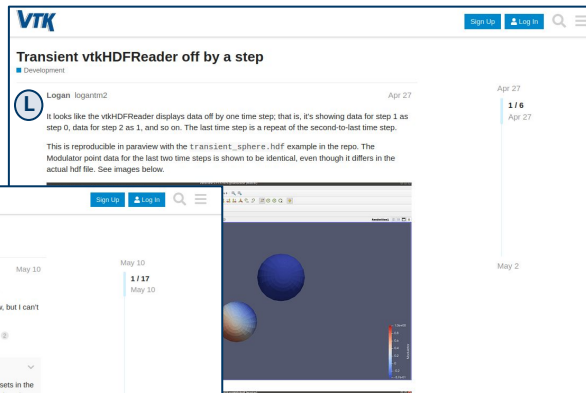
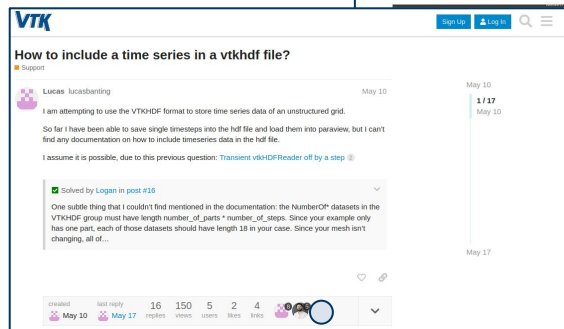


# ParaView's development drivers and feedback loops:

## Case Study

## Real User Feedback

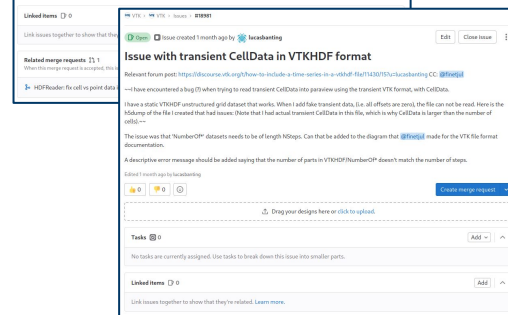
Forum User



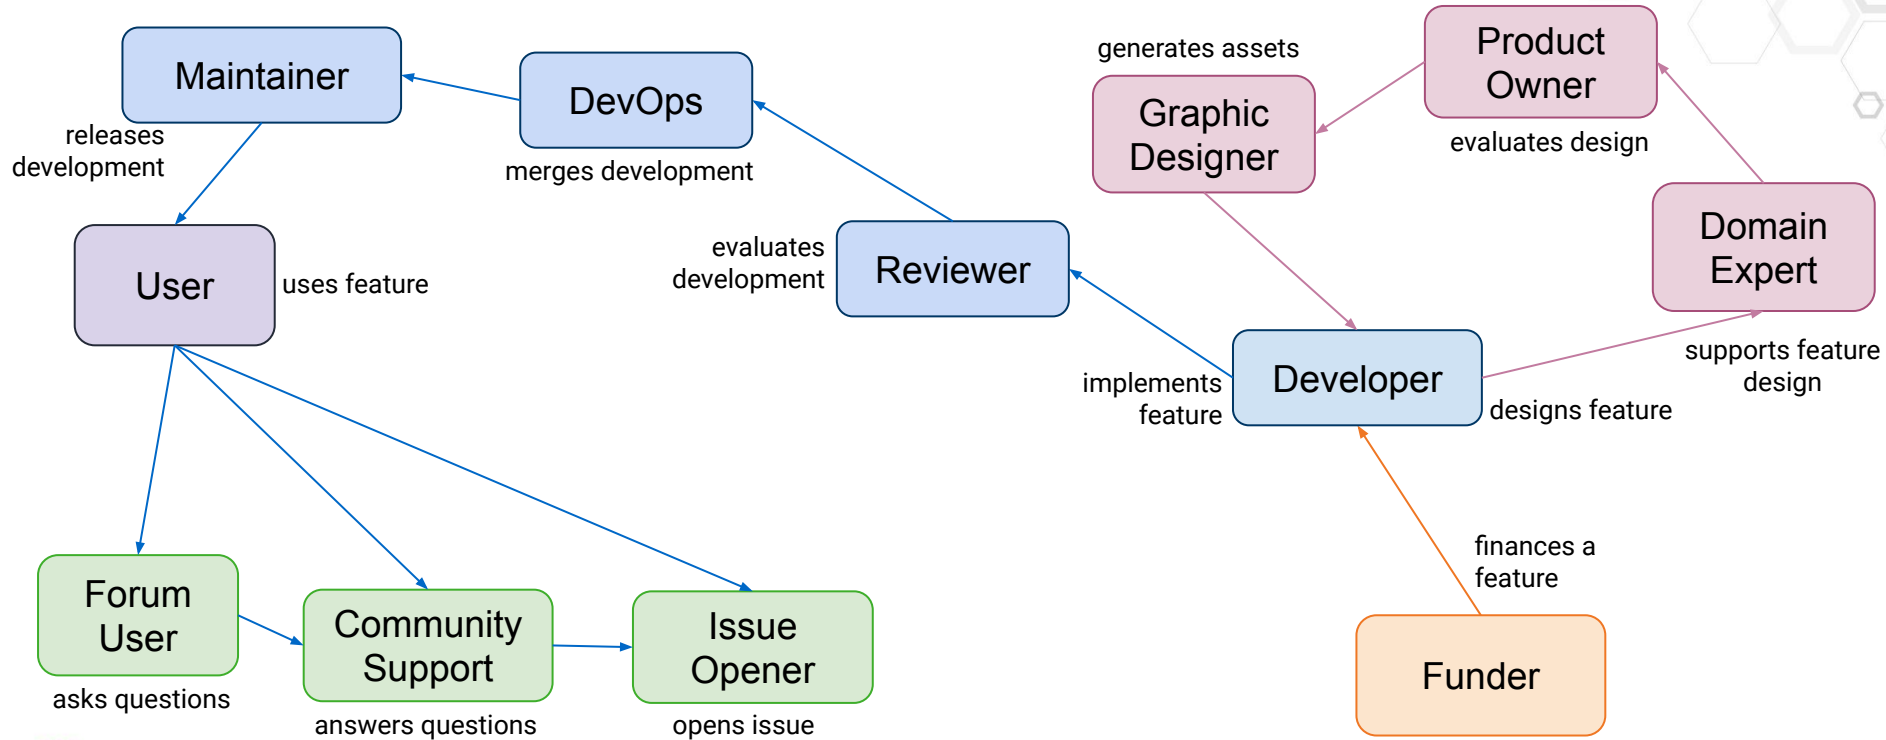
Community Support



Issue Opener



# ParaView's development drivers and feedback loops



# ParaView's development drivers and feedback loops: Case Study

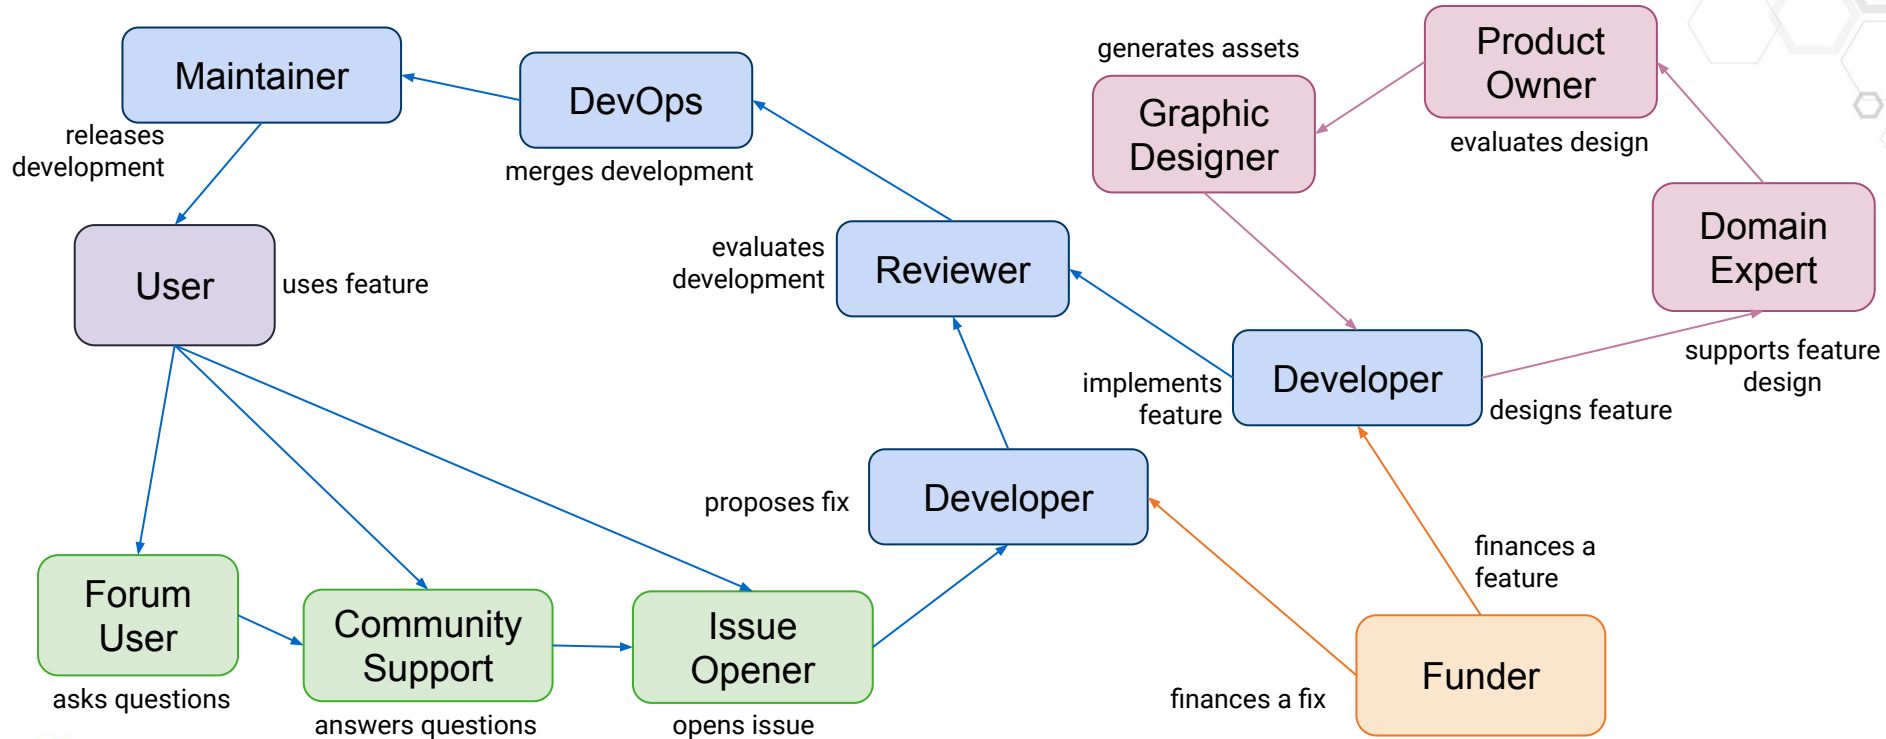
## Bug Fixing:

Developer

This screenshot shows a GitLab merge request titled "Fix transient image cell data" (ID #110238) requested by Julien Fausty. The request is merged and has been pushed to the master branch 3 weeks ago. The overview shows 10 commits, 2 pipelines, and 3 changes. The description states: "Fix the time dimension offsetting for all attributes except point data in the image data reading of vtkHDFReader." It also mentions "Fixes #18986 (closed)" and "Edited 3 weeks ago by Julien Fausty". The merge request is approved, with a note "Approval is optional". It was merged by Kitware Robot 2 weeks ago. The merge details include: "Changes merged into master with 7d6128a5", "Deleted the source branch", and "Closed #18986 (closed)".

This screenshot shows a GitLab merge request titled "vtkHDFReader: fix time step selection" (ID #110169) requested by Julien Fausty. The request is merged and has been pushed to the master branch 1 month ago. The overview shows 11 commits, 1 pipeline, and 1 change. The description states: "Consider that the time step of the data set below or equal to the requested time step is the time step to show" and "Related to: <https://discourse.vtk.org/t/transient-vtkhdfreader-off-by-a-step/11332>". It also mentions "Edited 1 month ago by Julien Fausty". The merge request is approved, with a note "Approval is optional". It was merged by Kitware Robot 1 month ago. The merge details include: "Changes merged into master with ed25ffa3" and "Deleted the source branch". The pipeline #330542 failed for d2ed2c2d 1 month ago. The pipeline #331023 was canceled for ed25ffa3 on master 1 month ago. The merge request was deployed to rsync-upload-wheel-sdk 1 month ago and to rsync-upload 1 month ago.

# ParaView's development drivers and feedback loops

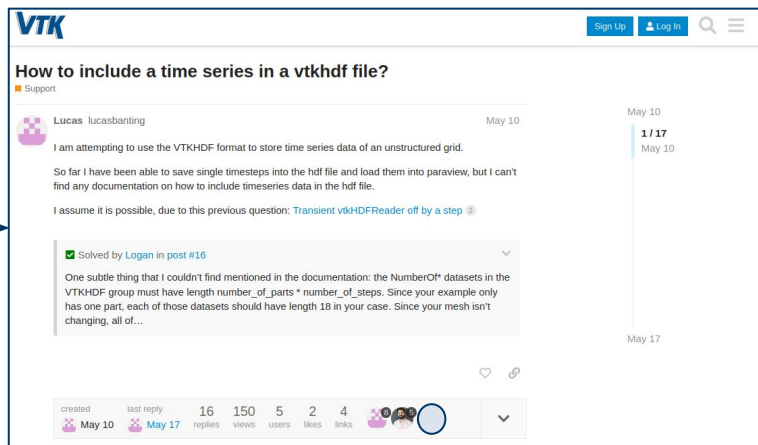


# Role transitioning in the ParaView eco-system

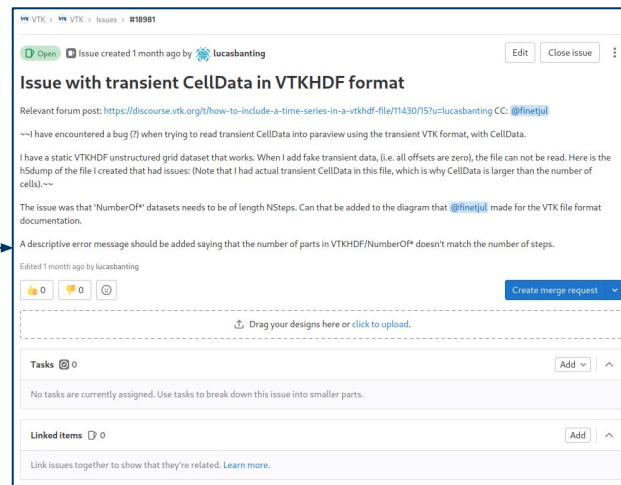
Theory and Practice

# Role transitioning in the ParaView eco-system: Case-study

User



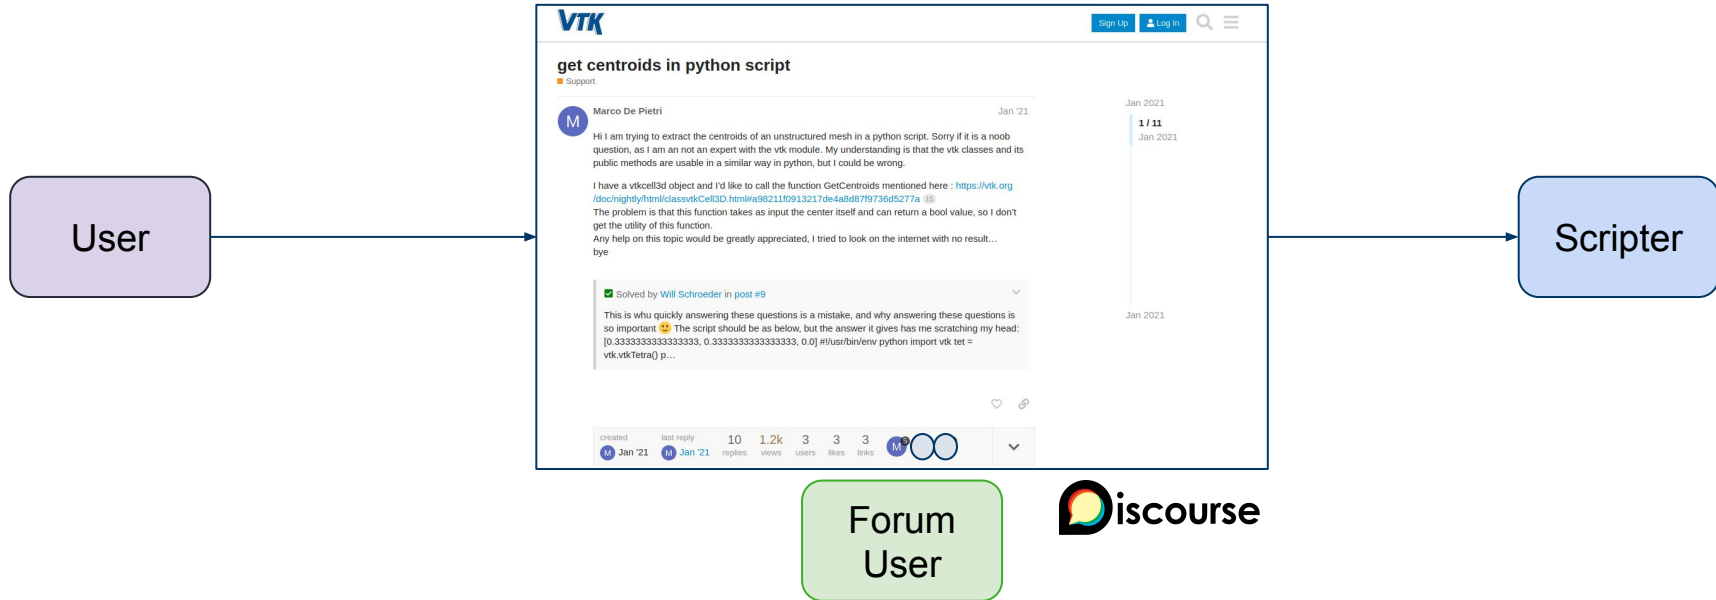
Forum  
User



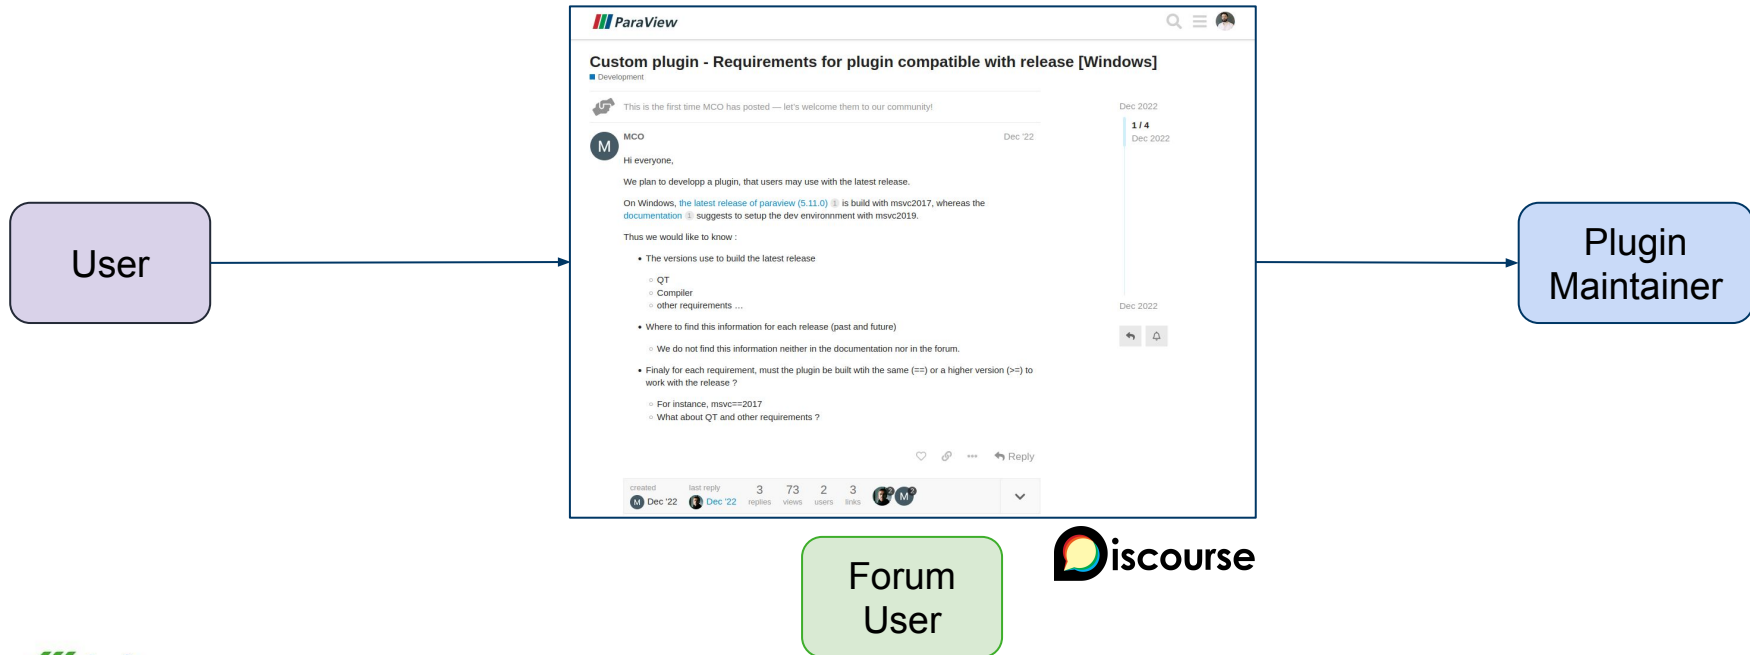
Issue  
Opener



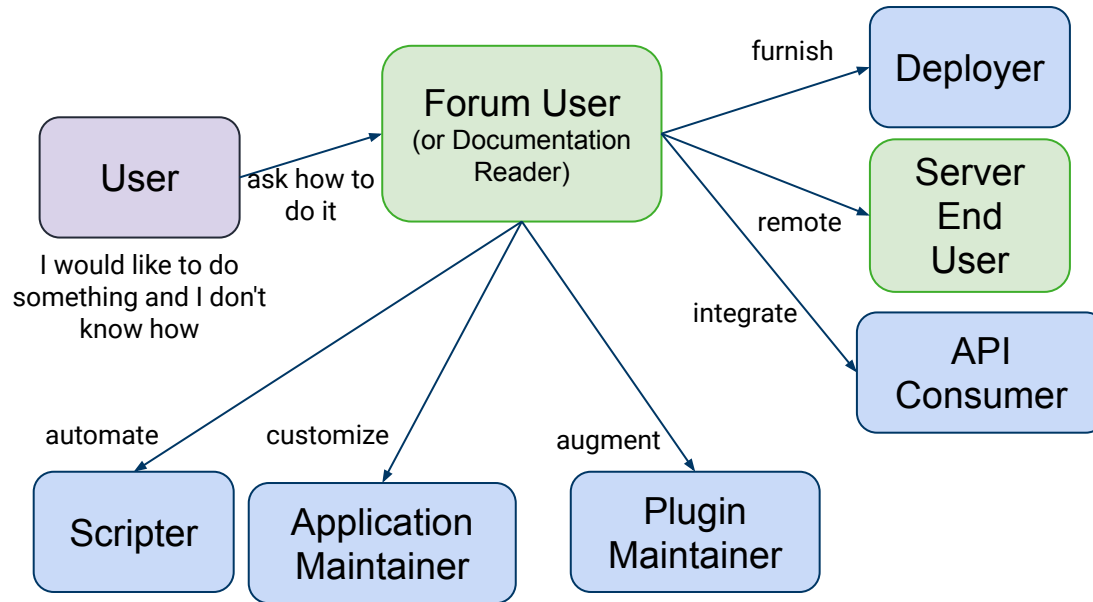
# Role transitioning in the ParaView eco-system: Case-study



# Role transitioning in the ParaView eco-system: Case-study

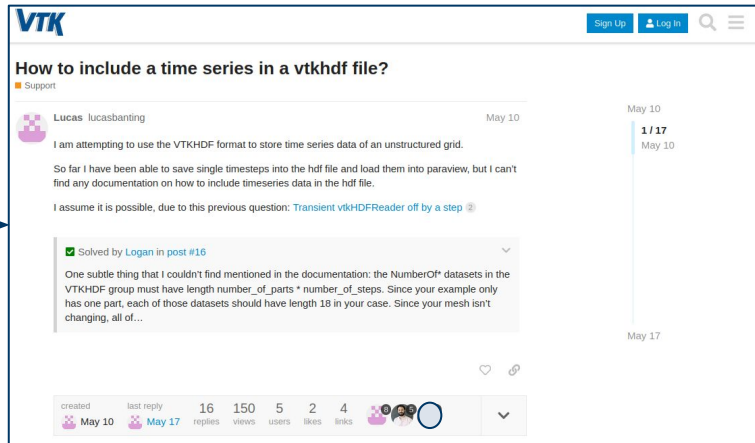


# Role transitioning in the ParaView eco-system: User-upskilling

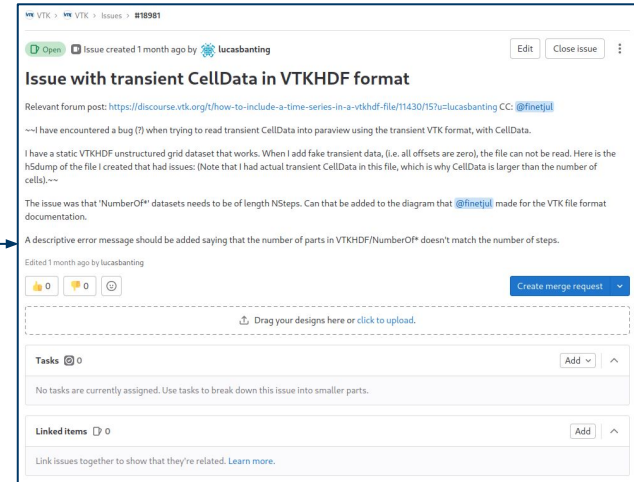


# Role transitioning in the ParaView eco-system: Case-study

User



Forum  
User

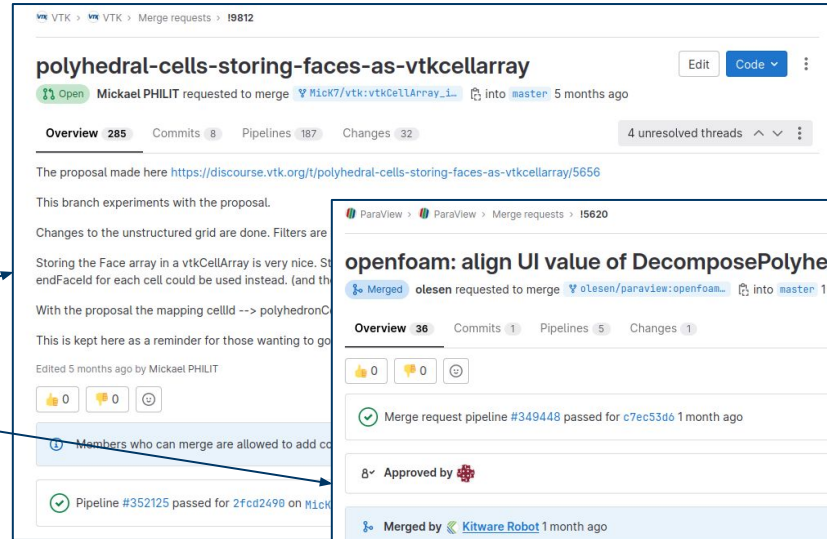


Issue  
Opener



# Role transitioning in the ParaView eco-system: Case-study

User



polyhedral-cells-storing-faces-as-vtkcellarray

Mickael PHILIT requested to merge `Mick7/vtk:vtkCellArray_1_` into `master` 5 months ago

Overview 285 Commits 8 Pipelines 187 Changes 32 4 unresolved threads

The proposal made here <https://discourse.vtk.org/t/polyhedral-cells-storing-faces-as-vtkcellarray/5656>

This branch experiments with the proposal.

Changes to the unstructured grid are done. Filters are

Storing the Face array in a `vtkCellArray` is very nice. Storing the `endFaceld` for each cell could be used instead. (and the

With the proposal the mapping `cellId --> polyhedronC`

This is kept here as a reminder for those wanting to go

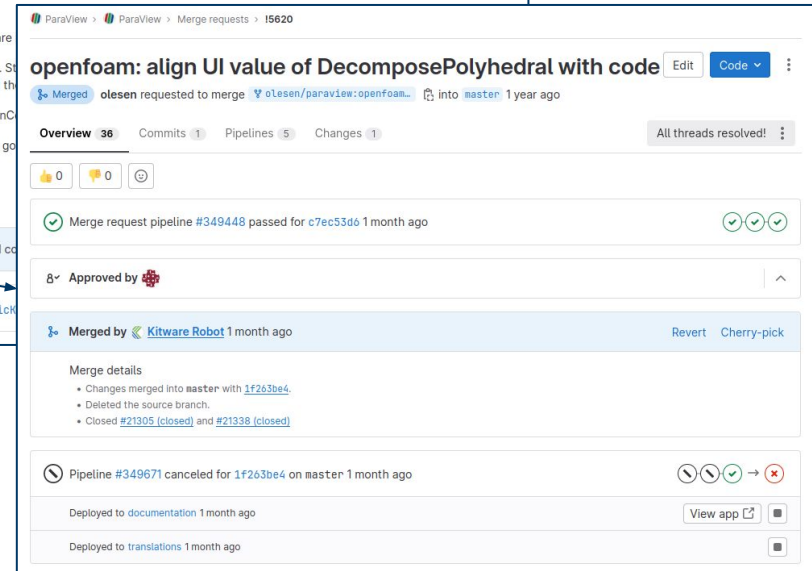
Edited 5 months ago by Mickael PHILIT

0 0 0

Members who can merge are allowed to add co

Pipeline #352125 passed for 2fcd2498 on h1ck

Developer



openfoam: align UI value of DecomposePolyhedral with code

olesen requested to merge `olesen/paraview:openfoam_` into `master` 1 year ago

Merged

Overview 36 Commits 1 Pipelines 5 Changes 1 All threads resolved!

0 0 0

Merge request pipeline #349448 passed for c7ec53d6 1 month ago

Approved by

Merged by Kitware Robot 1 month ago Revert Cherry-pick

Merge details

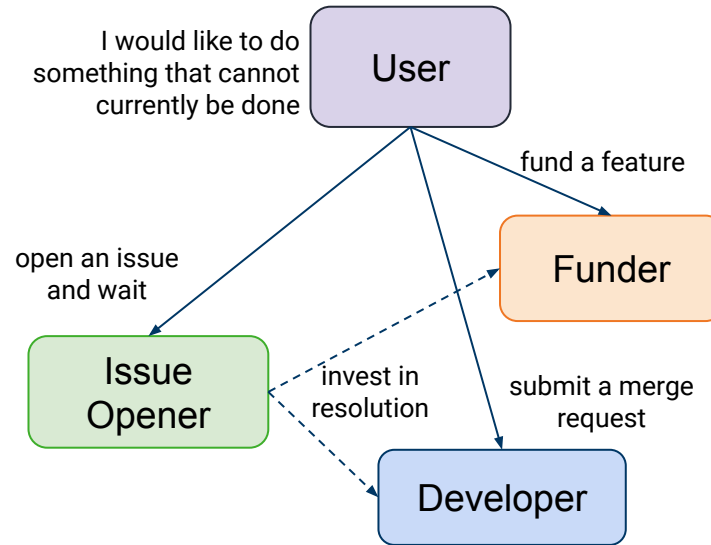
- Changes merged into `master` with `1f263be4`.
- Deleted the source branch.
- Closed [#21305](#) (closed) and [#21338](#) (closed)

Pipeline #349671 canceled for 1f263be4 on master 1 month ago

Deployed to documentation 1 month ago View app

Deployed to translations 1 month ago

# Role transitioning in the ParaView eco-system: User-Contributor transition



# Role transitioning in the ParaView eco-system: Case-study

Developer

Reviewer

ParaView > ParaView > Merge requests > 16306

### CleanToGrid: expose point data reduction strategies

Merged by Julien Fausty requested to merge v Julien.fausty/paraview:8... into master 1 month ago

Overview 5 Commits 3 Pipelines 2 Changes 8

The goal of this MR is to expose a new property in the "Clean to Grid" filter allowing users to choose a point data reduction strategy when points are merged during the clean operation.

0 0 0

✓ Merge request pipeline #330161 passed for 5c09e80c 1 month ago

8- Approval is optional

Merged by Kitware Robot 1 month ago Revert Cherry-pick

Merge details

- Changes merged into master with c4495c8c.
- Deleted the source branch.

✓ Pipeline #330281 passed for c4495c8c on master 1 month ago

Deployed to documentation 1 month ago View app

Deployed to translations 1 month ago

ParaView > ParaView > Merge requests > 16347

### Add zSpace plugin to CI

Merged by Thomas Galland requested to merge v thomas.galland/paraview:3... into master 3 months ago

Overview 15 Commits 2 Pipelines 6 Changes 4

The zSpace ParaView plugin does not need to search for zSpace Core compatibility headers and libraries. ZSPACE\_USE\_COMPAT\_SDK is set to ON, which is by default. Since this plugin does not have any compatibility headers, it is not necessary to build the CI to test that it builds correctly.

Needs: vtk/vtk10227 (merged)

Edited 1 week ago by Thomas Galland

0 0 0

✓ Merge request pipeline #353149 passed for 2a0d7899 6 days ago

8- Approval is optional

Merged by Kitware Robot 6 days ago Revert Cherry-pick

Merge details

- Changes merged into master with 4a0b790b.
- Deleted the source branch.

ParaView > ParaView > Merge requests > 16298

### vtkLogoSourceRepresentation: Expose border params

Merged by Charly Bollinger requested to merge v charly.bollinger/paraview:2... into master 2 months ago

Overview 52 Commits 4 Pipelines 8 Changes 10 All threads resolved!

- Expose how to render the border (always, never, on hover).
- Expose border thickness and color.
- Add a property to set scale (height relative to screen size).
- Fixed default scale so that borders fit the image.

0 0 0

✓ Merge request pipeline #329437 passed for 6c928d95 1 month ago

8- Approval is optional

Merged by Kitware Robot 1 month ago Revert Cherry-pick

Merge details

- Changes merged into master with 6b0332cf.
- Deleted the source branch.



Thomas Galland @thomas.galland · 1 month ago

+2 pending CI

Developer

Charly Bollinger @charly.bollinger · 1 month ago

+2

Developer

Julien Fausty @julien.fausty · 1 week ago

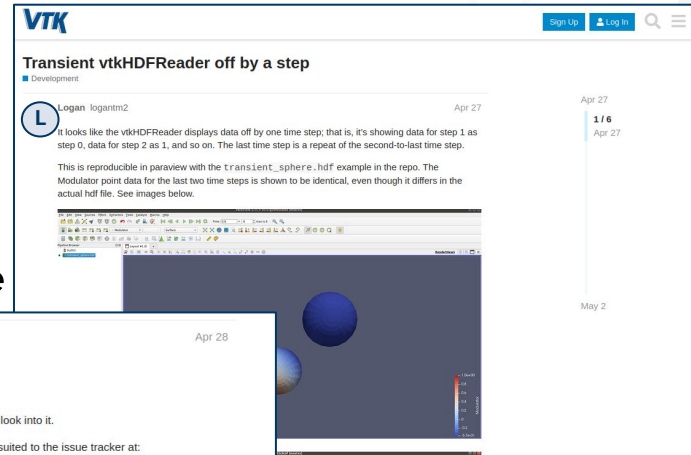
+2

Developer

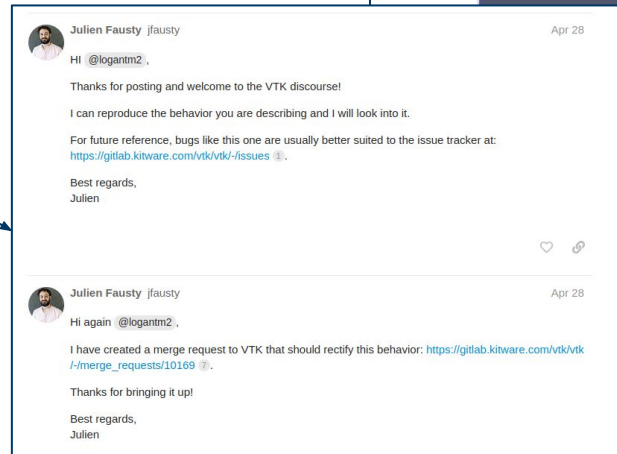
Edited 1 week ago by Julien Fausty

# Role transitioning in the ParaView eco-system: Case-study

Developer

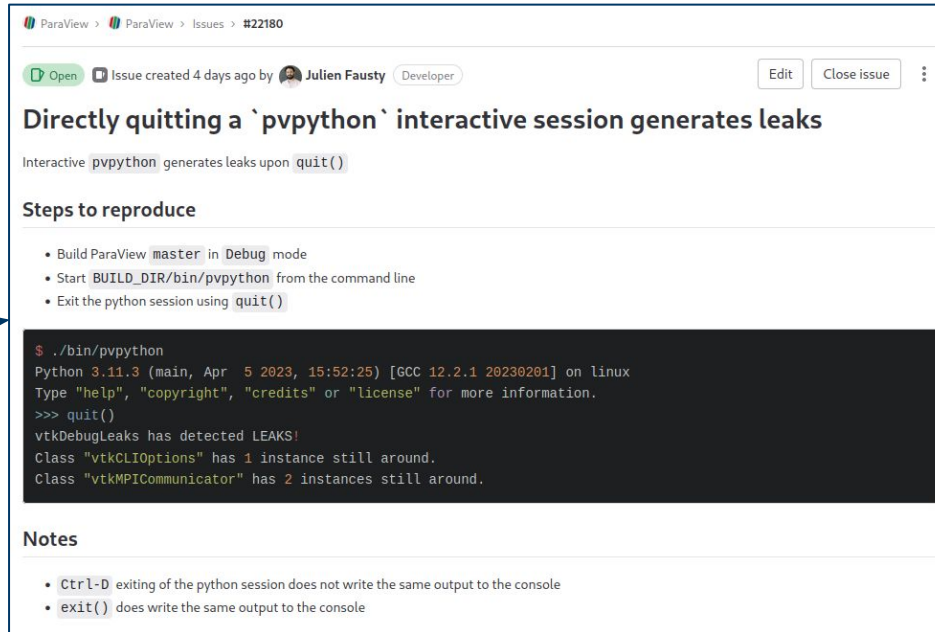


Community  
Support



# Role transitioning in the ParaView eco-system: Case-study

Developer



ParaView > ParaView > Issues > #22180

Open Issue created 4 days ago by **Julien Fausty** (Developer) Edit Close issue

## Directly quitting a `pvpython` interactive session generates leaks

Interactive `pvpython` generates leaks upon `quit()`

### Steps to reproduce

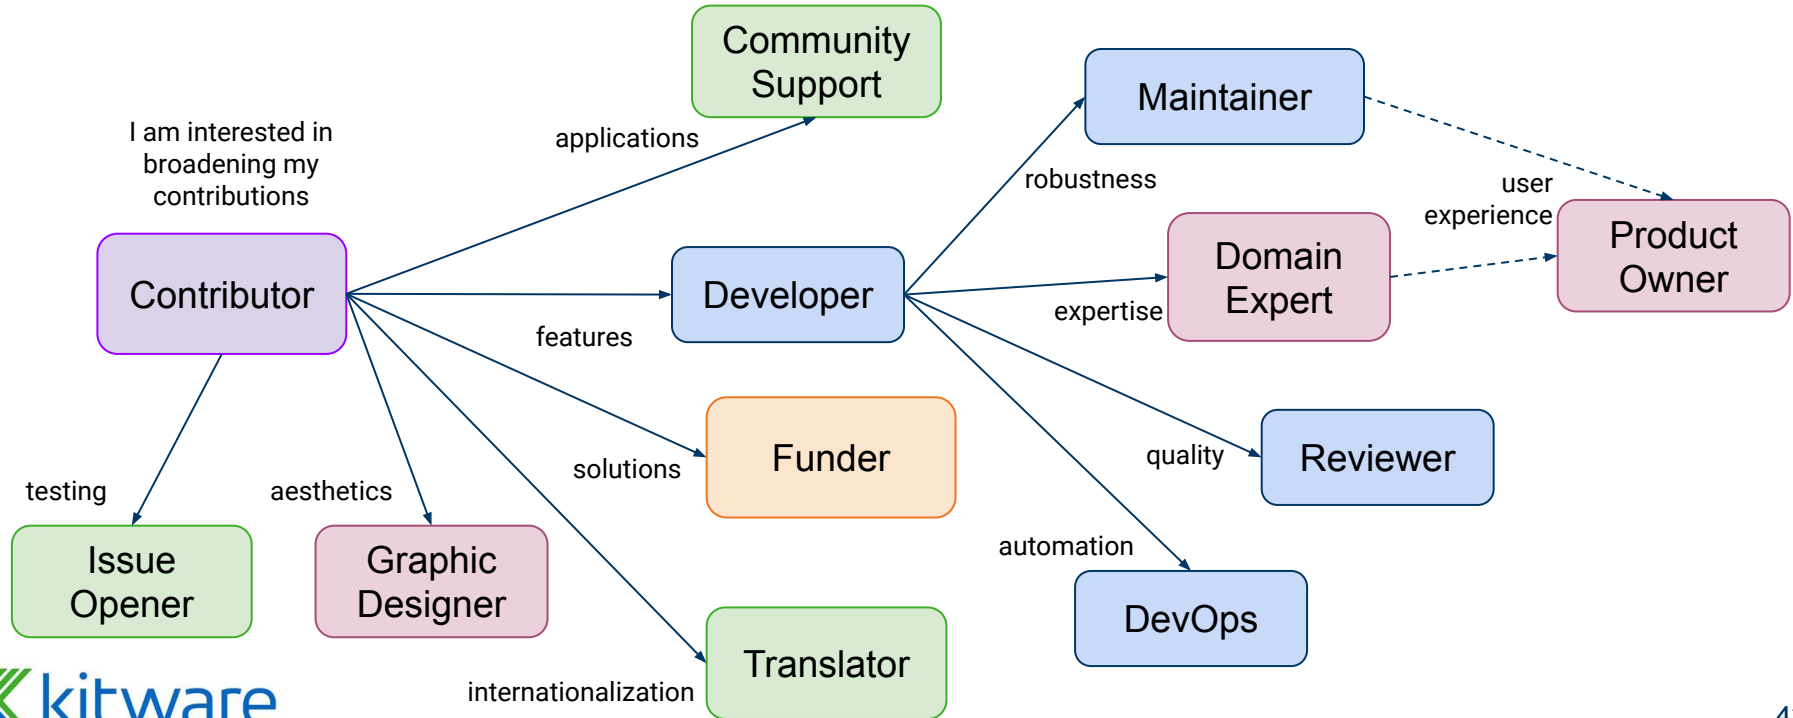
- Build ParaView `master` in `Debug` mode
- Start `BUILD_DIR/bin/pvpython` from the command line
- Exit the python session using `quit()`

```
$ ./bin/pvpython
Python 3.11.3 (main, Apr 5 2023, 15:52:25) [GCC 12.2.1 20230201] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> quit()
vtkDebugLeaks has detected LEAKS!
Class "vtkCLIOptions" has 1 instance still around.
Class "vtkMPICommunicator" has 2 instances still around.
```

### Notes

- `Ctrl-D` exiting of the python session does not write the same output to the console
- `exit()` does write the same output to the console

# Role transitioning in the ParaView eco-system: Contributor-upskilling



# The HPC Scientific Community

and final comments

# The particularities of serving a community of scientists

- **Innovative:** always pushing the boundaries of the tools it uses.
- **Curious:** to know how things work in the background.
- **Open:** to technical discussions and collaborations.
- **Diverse:** a variety of types of problems as well as cultures.

*Fosters a ParaView community with the following values:*

- **Communication:** the needs of the community is the driver for change
- **Open:** all development happens in the public domain favoring rigor and honesty
- **Efficient:** every development resource counts

- **DIY:** likely to use home grown solutions instead of industry standards.
- **KPIs:** indicators are skewed towards final results and not tooling.
- **Time-Crunched:** many interesting things to do and not enough time to do them all.

# The challenges of HPC open-source tools

Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
1	<b>Frontier</b> - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 26GHz, AMD Instinct MI250X, Slingshot-11, HPE DOE/SC/Oak Ridge National Laboratory United States	8,699,904	1,194.00	1,679.82	22,703
2	<b>Supercomputer Fugaku</b> - Supercomputer Fugaku, A64FX 48C 2.26GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442.01	537.21	29,899
3	<b>LUMI</b> - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 26GHz, AMD Instinct MI250X, Slingshot-11, HPE EuroHPC/CSC Finland	2,220,288	309.10	428.70	6,016
4	<b>Leonardo</b> - BullSequana XH2000, Xeon Platinum 8358 32C 2.6GHz, NVIDIA A100 SXM4 64 GB, Quad-rail NVIDIA HDR100 Infiniband, Atos EuroHPC/CINECA Italy	1,824,768	238.70	304.47	7,404
5	<b>Summit</b> - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148.60	200.79	10,096

~ 600 million \$

[https://en.wikipedia.org/wiki/Frontier\\_\(supercomputer\)](https://en.wikipedia.org/wiki/Frontier_(supercomputer))

~ 1 billion \$

(total project cost)

[https://en.wikipedia.org/wiki/Fugaku\\_\(supercomputer\)](https://en.wikipedia.org/wiki/Fugaku_(supercomputer))

~ 144.5 million €

<https://en.wikipedia.org/wiki/LUMI>

~ 240 million €

[https://en.wikipedia.org/wiki/Leonardo\\_\(supercomputer\)](https://en.wikipedia.org/wiki/Leonardo_(supercomputer))

~ 325 million \$

(with Sierra #6)

[https://en.wikipedia.org/wiki/Summit\\_\(supercomputer\)](https://en.wikipedia.org/wiki/Summit_(supercomputer))

- **Access:** target architectures are more difficult / expensive to access and test on.
- **Heterogeneity:** no "one size fits all" for HPC machines or even within a single HPC machine with heterogeneous hardware.

**Solution:** rely on **users** to test in HPC settings and collaborate to improve support.

# Some conclusions and outlook

Today, ParaView owes much of its success to its community being:

- **Diverse:** academic and industrial on an international scale
- **Innovative:** constantly pushing the boundaries
- **Open:** to transitioning roles and becoming more active
- **User Oriented:** seeing as every member is also a user.

## Outlook

- Opportunities around adapting software for the exascale
- Recognition of open-source projects as rigorous frameworks for development
- Good synergy between reproducible science and open-source frameworks with large communities



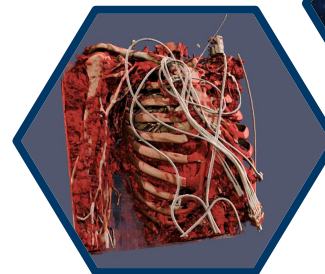
# Take Part in the First ParaView User Day Europe!

## Lyon, France - September 19th, 2023

- Inspiring keynotes
- Lightning talks by users
- Kitware one-to-one hands-on sessions
- Demo space
- ParaView dinner (optional)

Join the experience and present!

We understand the value of your expertise, if you want to share with the community, you can submit a presentation title and summary to present during the event.



**Thank you for your attention and looking forward to your contributions ;) !**

Questions?

Or/and contact us at [kitware@kitware.fr](mailto:kitware@kitware.fr) or [julien.fausty@kitware.com](mailto:julien.fausty@kitware.com)