



CIPRES

Cyberinfrastructure for
Phylogenetic Research



The CIPRES Science Gateway: A Community Resource for Phylogenetic Analyses

Mark Miller, Wayne Pfeiffer, and Terri Schwartz

San Diego Supercomputer Center



POWERED BY  TeraGrid™

SDSC




CIPRES

Cyberinfrastructure for
Phylogenetic Research



Systematics is the study of biological diversity and its origins. It focuses on understanding evolutionary relationships among organisms.



?



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



Originally, evolutionary relationships were inferred from morphology alone:

Morphological characters are scored “by hand” to create matrices of characters.

Scoring occurs via low volume/low throughput methodologies

Even though tree inference is NP hard, matrices created using morphological characters alone are typically relatively small, so computations are relatively tractable (with heuristics developed by the community)



POWERED BY  TeraGrid™

SDSC




Evolutionary relationships are now inferred from DNA sequence comparisons:

Align sequences to determine evolutionary equivalence:

species	263	*	260	*	300	*	320		
species 1	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCCCAATTAAAG	GTGA	AACCBCAATGGCTCAITAAA	TCA			
species 2	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAATTCCTTTGA	GG	GAGA	AATTCBAAAGGCTCAITAAA	TCA		
species 3	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCTTCACTAAAG	GTGA	AACCBCAATGGCTCAITAAA	TCA			
species 4	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAGMCCG	ATCT	AAG	GCGA	AACCBCAATGGCTCAITAAA	TCA	
species 5	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAGGCCG	ATCT	AAG	GCGA	AACCBCAATGGCTCAITAAA	TCA	
species 6	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAGGCCG	AACT	AAG	GCGA	AACCBCAATGGCTCAITAAA	TCA	
species 7	TGGTFTCTCGTTT	CCCTGC	TGCTAAAGT	ACAAGCCG	ATTG	AAG	GCGA	AACCBCAATGGCTCAITAAA	TCA
species 8	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCCG	ATTG	AAG	GCGA	AACCBCAATGGCTCAITAAA	TCA	
species 9	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCCG	ATCT	AAG	GCGA	AACCBCAATGGCTCAITAAA	TCA	
species 10	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACA	---	CCTCTG	GG	GCGA	AACCBCAATGGCTCAITAAA	TCA
species 11	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCCG	TATG	CG	GCGA	AACCBCAATGGCTCAITAAA	TCA	
species 12	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCCG	TAGA	CG	GCGA	AACCBCAATGGCTCAITAAA	TCA	
species 13	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ATAAGCCCAATA	AA	---	GTGA	GACCBCAATGGCTCAITAAA	TCA	
species 14	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCCG	CAATA	AAT	GTGA	GACCBCAATGGCTCAITAAA	TCA	
species 15	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCCG	TATA	TATG	GTGA	GACCBCAATGGCTCAITAAA	TCA	
species 16	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCCG	CAATA	AA	---	GTGA	GACCBCAATGGCTCAITAAA	TCA
species 17	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCCG	TACA	GG	---	GTGA	AACCBCAATGGCTCAITAAA	TCA
species 18	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCCG	CAATA	AA	---	GTGA	AACCBCAATGGCTCAITAAA	TCA
species 19	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCCG	CAATA	AA	---	GTGA	AACCBCAATGGCTCAITAAA	TCA
species 20	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCCG	TACA	CG	---	GTGA	AACCBCAATGGCTCAITAAA	TCA
species 21	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCCG	TACA	CG	---	GTGA	AACCBCAATGGCTCAITAAA	TCA
species 22	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCCG	TACA	CG	---	GTGA	AACCBCAATGGCTCAITAAA	TCA
species 23	TCAAAGATTAAAGC	CATGCATGCTAAAGT	ACAAGCCG	TACA	CG	---	GTGA	AACCBCAATGGCTCAITAAA	TCA

Infer evolutionary relationships based on some set of assumptions:





CIPRES

Cyberinfrastructure for
Phylogenetic Research



Inferring Evolutionary relationships from DNA sequence comparisons is powerful:

DNA sequences are determined by fully automated procedures.

Sequence data can be gathered from many species at scales from gene to whole genome.

The high speed and low cost of NexGen Sequencing means new levels of sensitivity and resolution can be obtained.



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



Inferring Evolutionary relationships from DNA sequence comparisons is powerful, **BUT:**

Current analyses often involve 1000's of species and 1000's of characters, creating very large matrices.

Sequence alignment and Tree inference are NP hard, so even with heuristics, computational power often limits the analyses (already).

The length of tree search analysis scales exponentially with number of taxa and with number of characters with codes in current use.

There are at least 10^7 species, each with 3000 - 30,000 genes, so the need for computational power and new approaches will continue to grow.



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



Inferring Evolutionary relationships from DNA sequence comparisons is powerful, **BUT:**

Current analyses often involve **1000's of species and 1000's of characters**, creating very large matrices.

Sequence alignment and Tree inference are NP hard, so even with heuristics, computational power often limits the analyses (already).

The length of tree search analysis scales exponentially with number of taxa and with number of characters with codes in current use.

There are at least 10^7 species, each with 3000 - 30,000 genes, so the need for computational power and new approaches will continue to grow.



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



Inferring Evolutionary relationships from DNA sequence comparisons is powerful, **BUT:**

Current analyses often involve **1000's of species and 1000's of characters**, creating very large matrices.

Sequence alignment and Tree inference are NP hard, so even with heuristics, **computational power often limits the analyses (already)**.

The length of tree search analysis scales exponentially with number of taxa and with number of characters with codes in current use.

There are at least 10^7 species, each with 3000 - 30,000 genes, so the need for computational power and new approaches will continue to grow.



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



Inferring Evolutionary relationships from DNA sequence comparisons is powerful, **BUT:**

Current analyses often involve **1000's of species and 1000's of characters**, creating very large matrices.

Sequence alignment and Tree inference are NP hard, so even with heuristics, **computational power often limits the analyses (already)**.

The length of tree search analysis **scales exponentially with number of taxa and with number of characters** with codes in current use.

There are at least 10^7 species, each with 3000 - 30,000 genes, so the need for computational power and new approaches will continue to grow.



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



Inferring Evolutionary relationships from DNA sequence comparisons is powerful, **BUT:**

Current analyses often involve **1000's of species and 1000's of characters**, creating very large matrices.

Sequence alignment and Tree inference are NP hard, so even with heuristics, **computational power often limits the analyses (already)**.

The length of tree search analysis **scales exponentially with number of taxa and with number of characters** with codes in current use.

There are at least **10^7 species, each with 3000 - 30,000 genes**, so the need for computational power and new approaches will continue to grow.



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



The CIPRES Project was created to support this new age of large phylogenetic data sets. The project had as its principal goals:

- 1. Developing heuristics and tools for analyzing the large DNA data sets that are available.**
- 2. Improving researcher access to computational resources.**



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



The CIPRES Portal was designed to allow users analyze large sequence data sets using community codes on a significant computational resource.

The CIPRES Portal provided:

- **Login-protected personal user space for storing results indefinitely.**
- **Access to most/all native command line options for each code.**
- **Support for adding new tools and new versions as needed.**

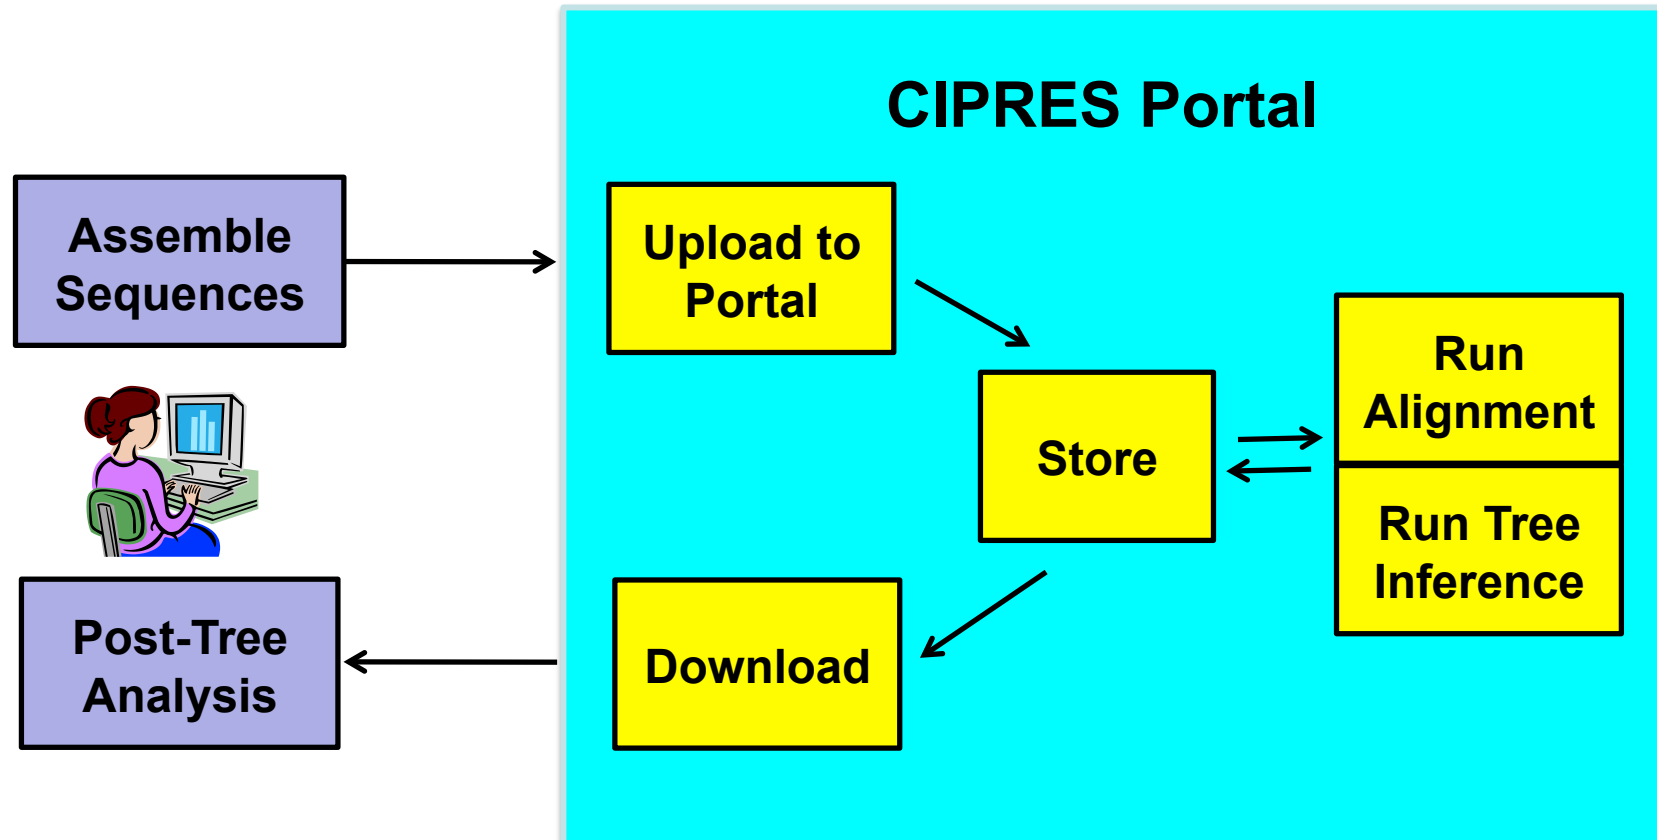


POWERED BY  TeraGrid™

SDSC

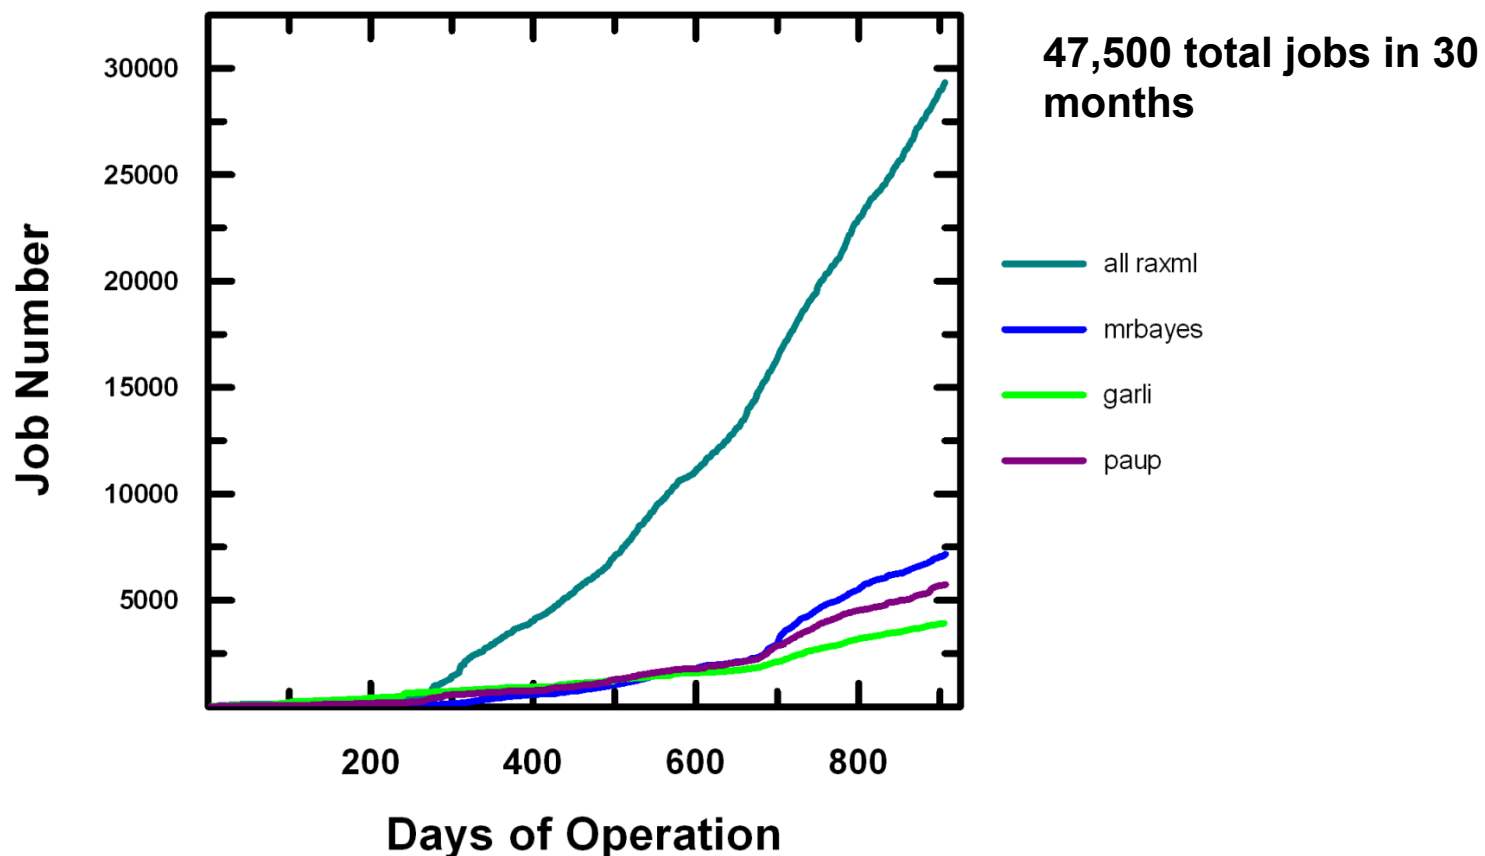


Workflow for the CIPRES Portal:





Usage Statistics for CIPRES Portal 5/2007 – 11/2009





CIPRES

Cyberinfrastructure for
Phylogenetic Research



Limitations of the original CIPRES Portal

- all jobs were run serially (efficient, but no gain in wall time)
- runs were limited to 72 hours
- demand for job runs was increasing
- the cluster was modest (16 X 8-way dual core nodes)
- the cluster was at the end of its useful lifetime
- funding for the project was ending

This is not a scalable, sustainable solution!



POWERED BY  TeraGrid

SDSC

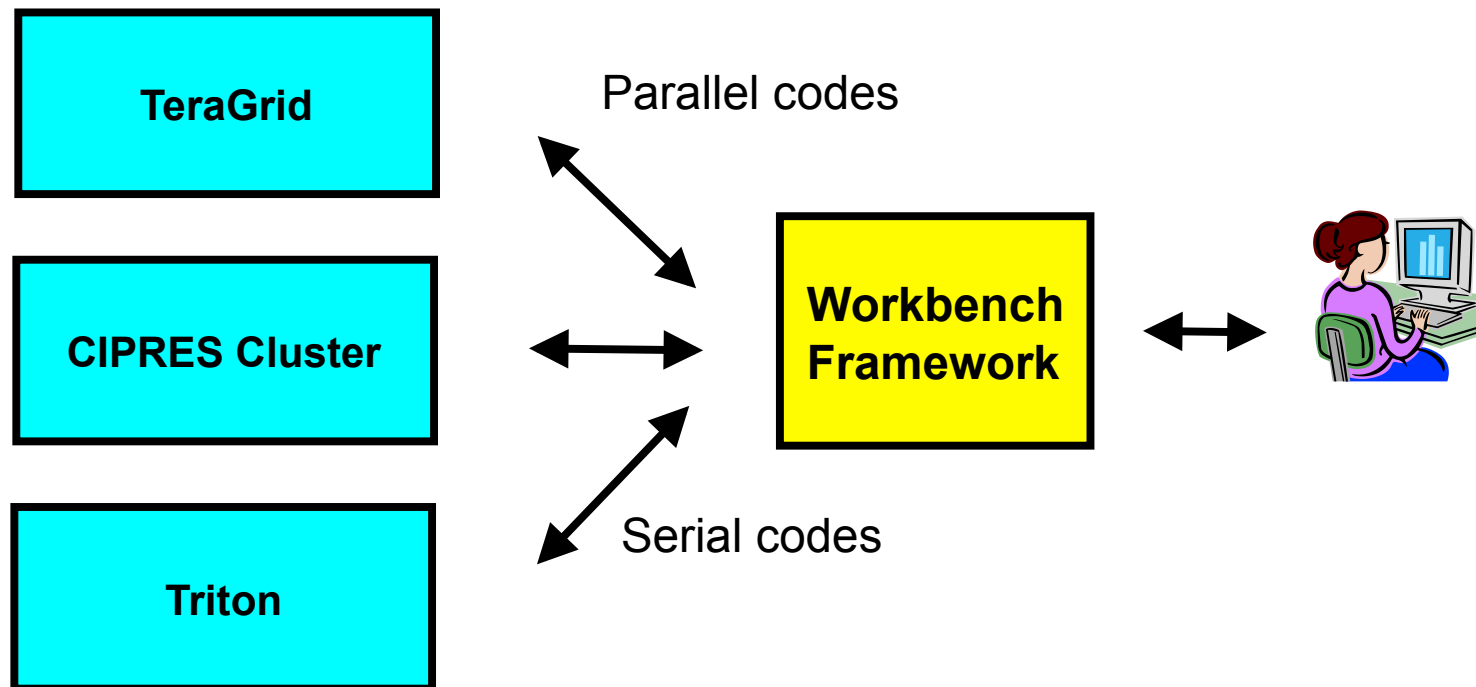


CIPRES

Cyberinfrastructure for
Phylogenetic Research



The solution: make community codes available on scalable, sustainable resources via Science Gateway Program

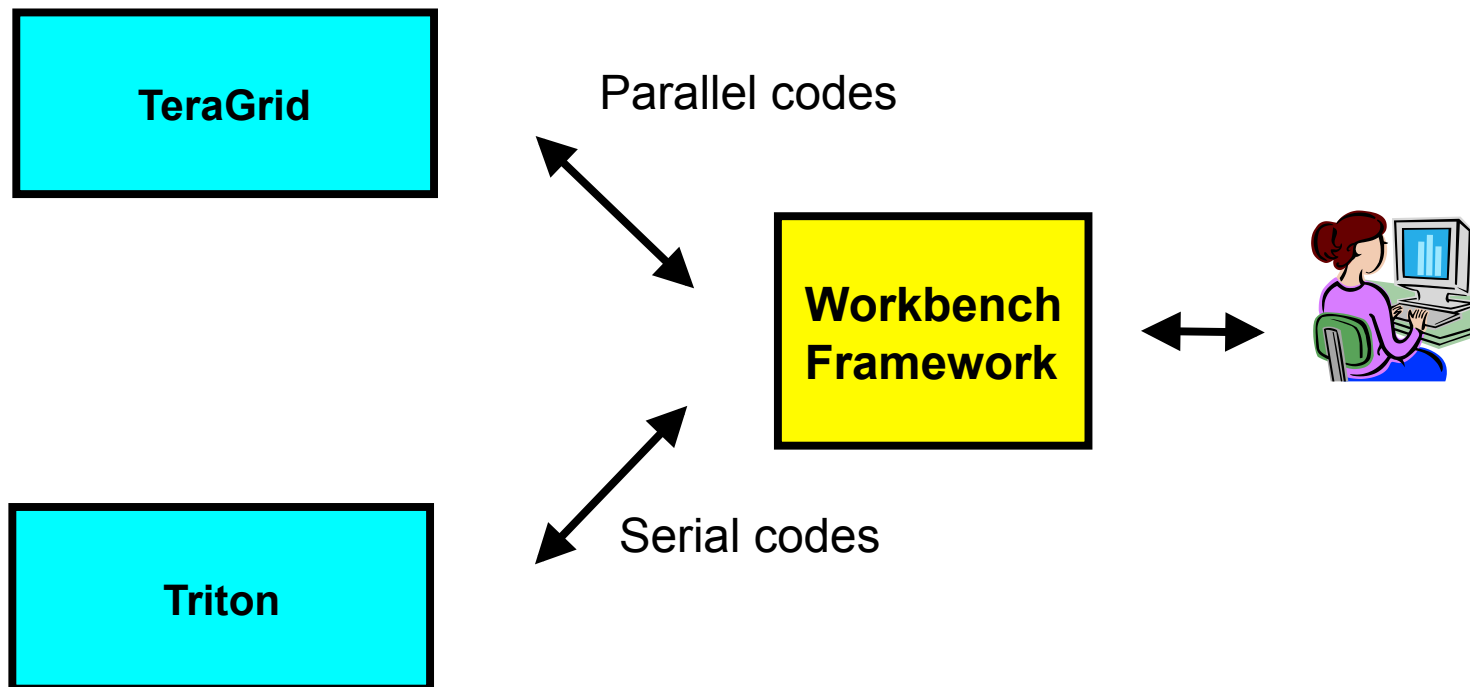


POWERED BY  TeraGrid™

SDSC



The solution: make community codes available on scalable, sustainable resources via Science Gateway Program





CIPRES

Cyberinfrastructure for
Phylogenetic Research



Greater than 90% of all computational time was used for three tree inference codes: MrBayes, RAxML, and GARLI.

Deploy parallel versions of these codes on TeraGrid Machines; initially using Globus/GRAM.

Work with community developers to improve the speed-up available through the parallel codes offered by CSG.

Add new parallel codes (e.g. MAFFT) as they appear in the community.

Keep other serial codes on local SDSC resources that provide the project with fee-for-service cycles.



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



CIPRES Science Gateway Code Improvements

- Hybrid MPI/Pthreads version of RAxML was developed
- Hybrid MPI/OpenMP version of MrBayes was developed
- Single-node runs are more efficient than before
- Multi-node runs with more cores are possible
 - Scalability before was limited to about 8 cores for typical analyses
 - Codes now scale well to 10s of cores for typical analyses



POWERED BY  TeraGrid

SDSC




Parallel code profiles on Trestles

Code	Type	Max cores	Speed-up	Efficiency
MrBayes	Hybrid MPI/OpenMP	32	2.4 X (1 nodes)	~60%
RAxML	Hybrid MPI/Pthreads	60	3.0 X (2 nodes)	~ 60%
GARLI	MPI	100	77 X (100 nodes)	77-94%
MAFFT	Pthreads	32	10 X	25%





What happens if you build it and too many people come???





CIPRES

Cyberinfrastructure for
Phylogenetic Research



What happens if you build it and too many people come???

- **make sure resource use is efficient**
- **make sure resource use delivers impact**
- **implement policies to maximize access for all**



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



What happens if you build it and too many people come???

- **make sure resource use is efficient**
- **make sure resource use delivers impact**
- **implement policies to maximize access for all**



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



Job Attrition on the CIPRES Science Gateway*



*March – August 2010



POWERED BY  TeraGrid™

SDSC



Error Impact analysis

	CPU time	User	Staff
Input error	0	med	low
Machine error	0	med	low
Communication error	high	high	high
Unknown error	?	med	low





CIPRES

Cyberinfrastructure for
Phylogenetic Research



CONCLUSION: Time to refactor the job monitoring system!

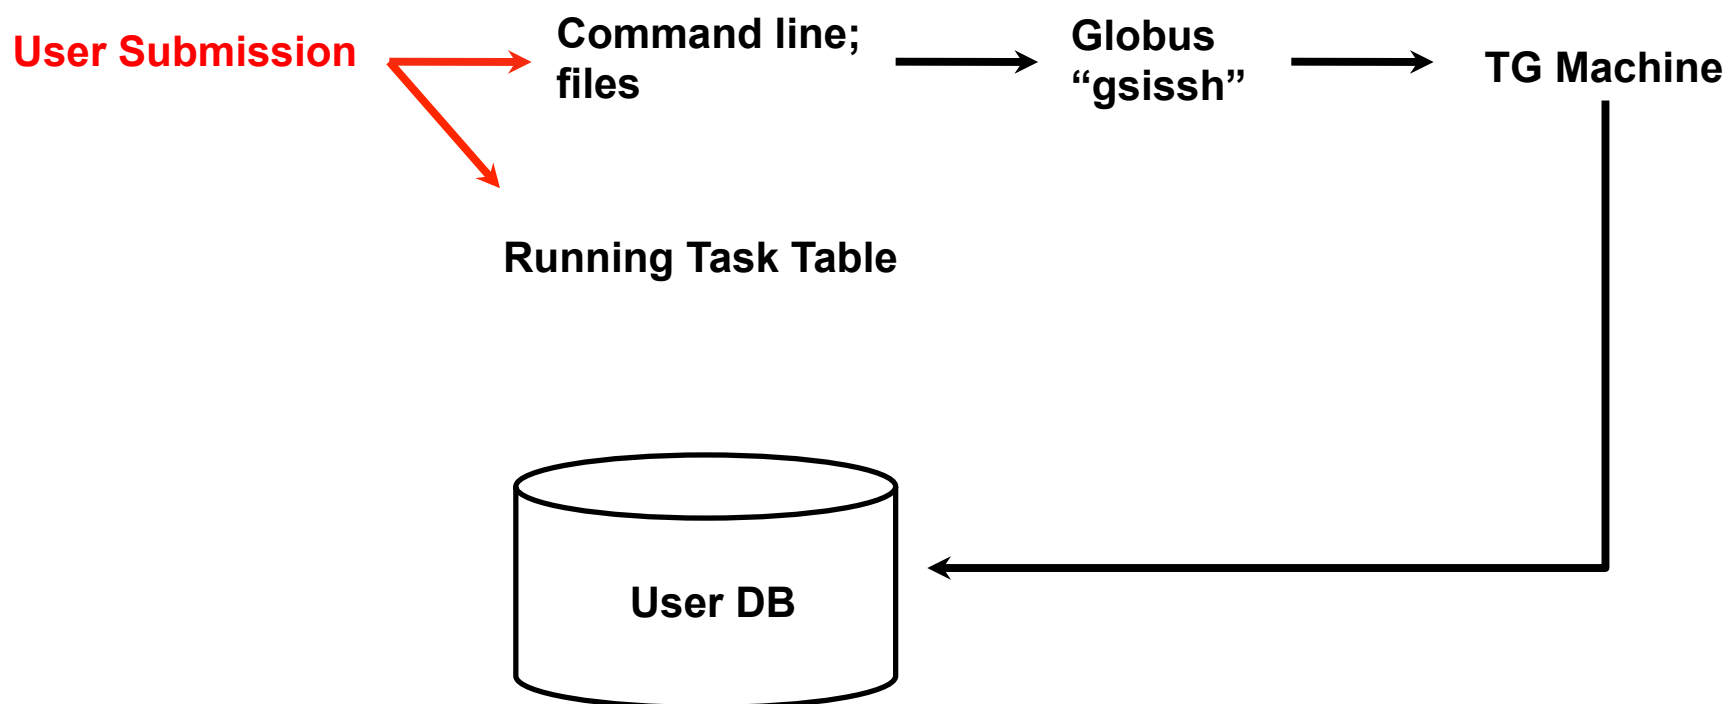


POWERED BY  TeraGrid™

SDSC

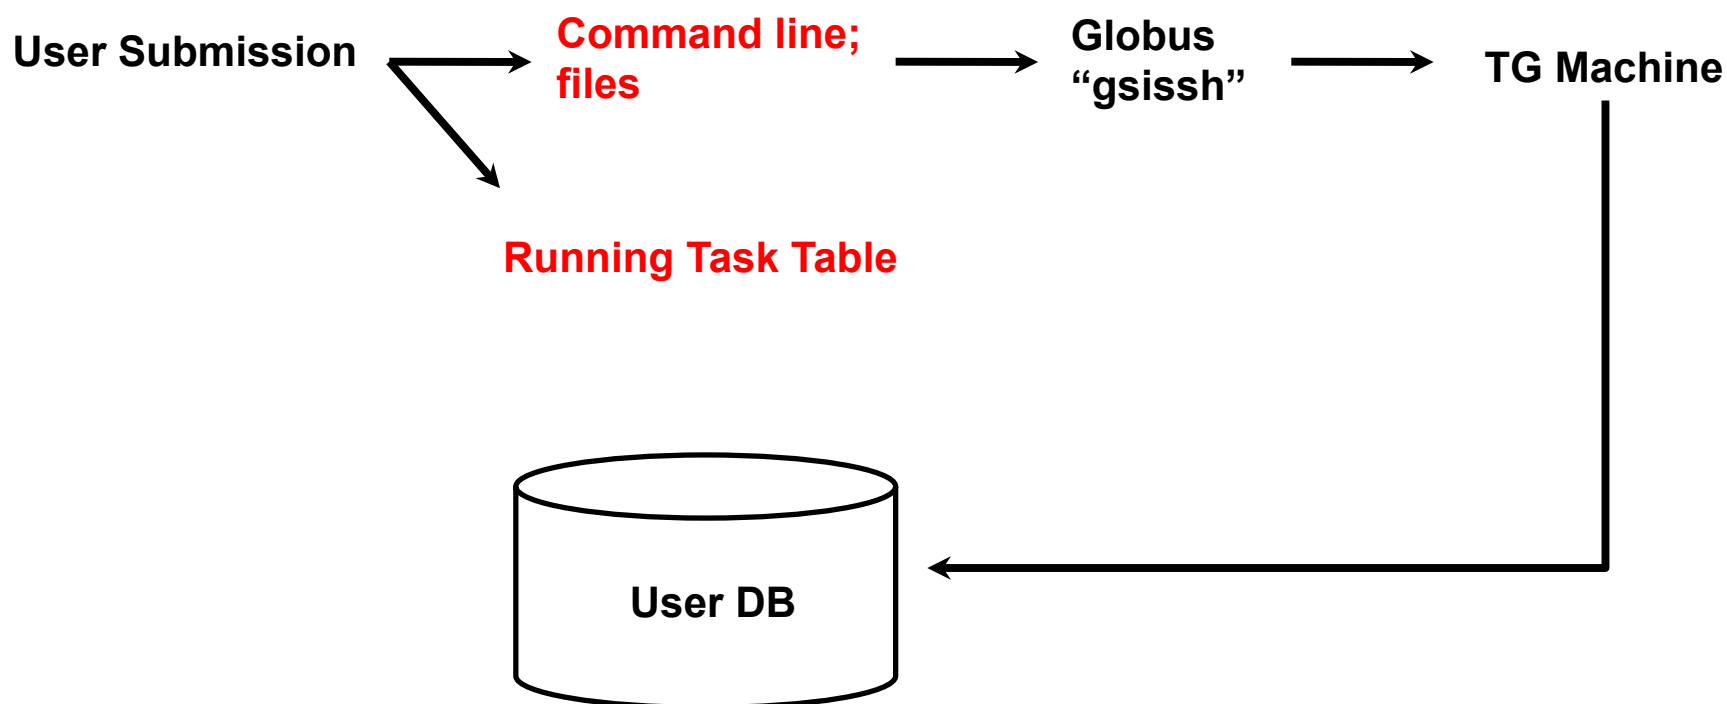


Normal Operation



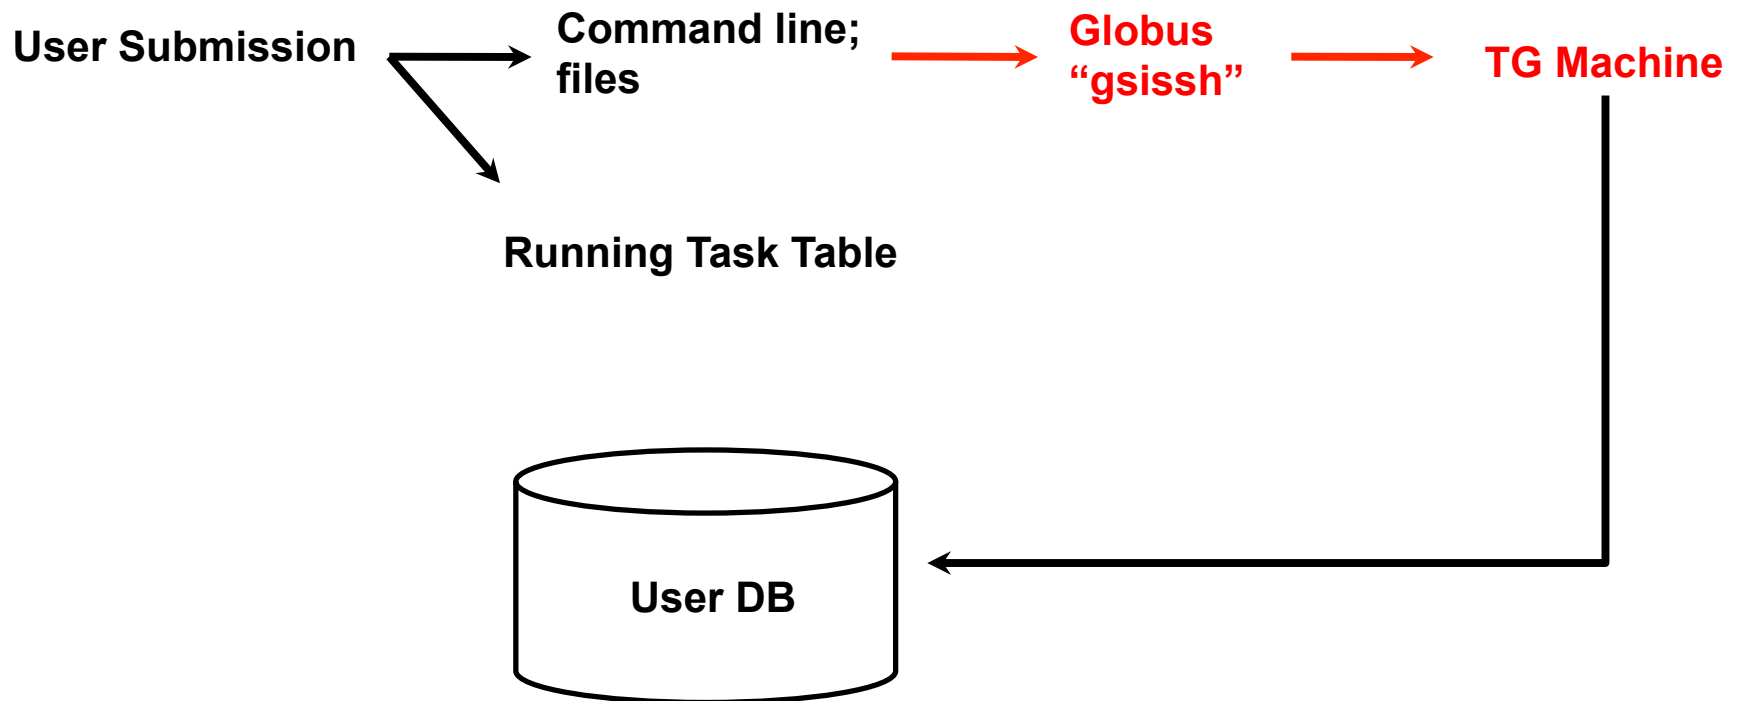


Normal Operation



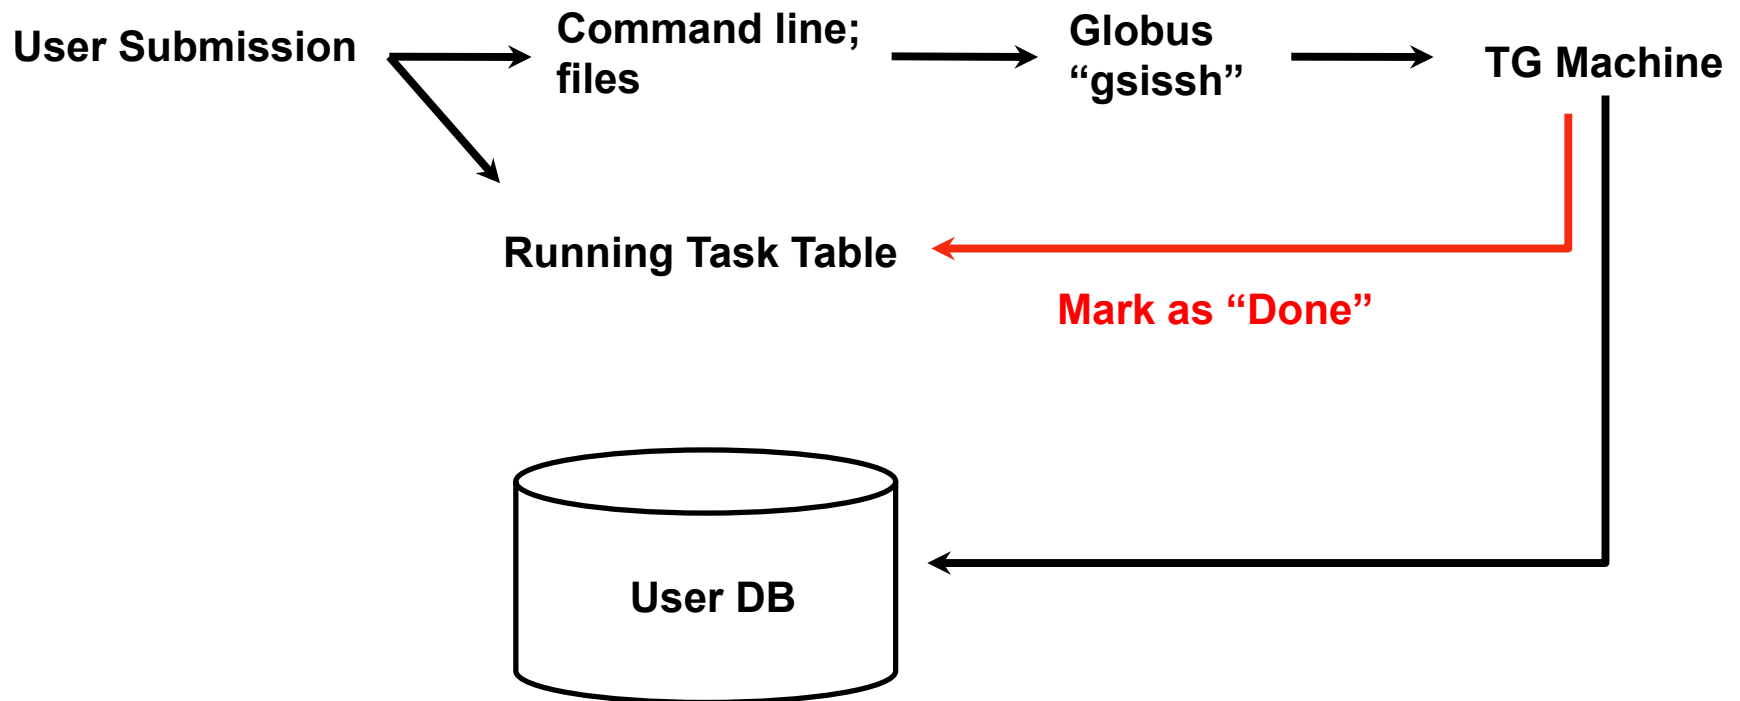


Normal Operation



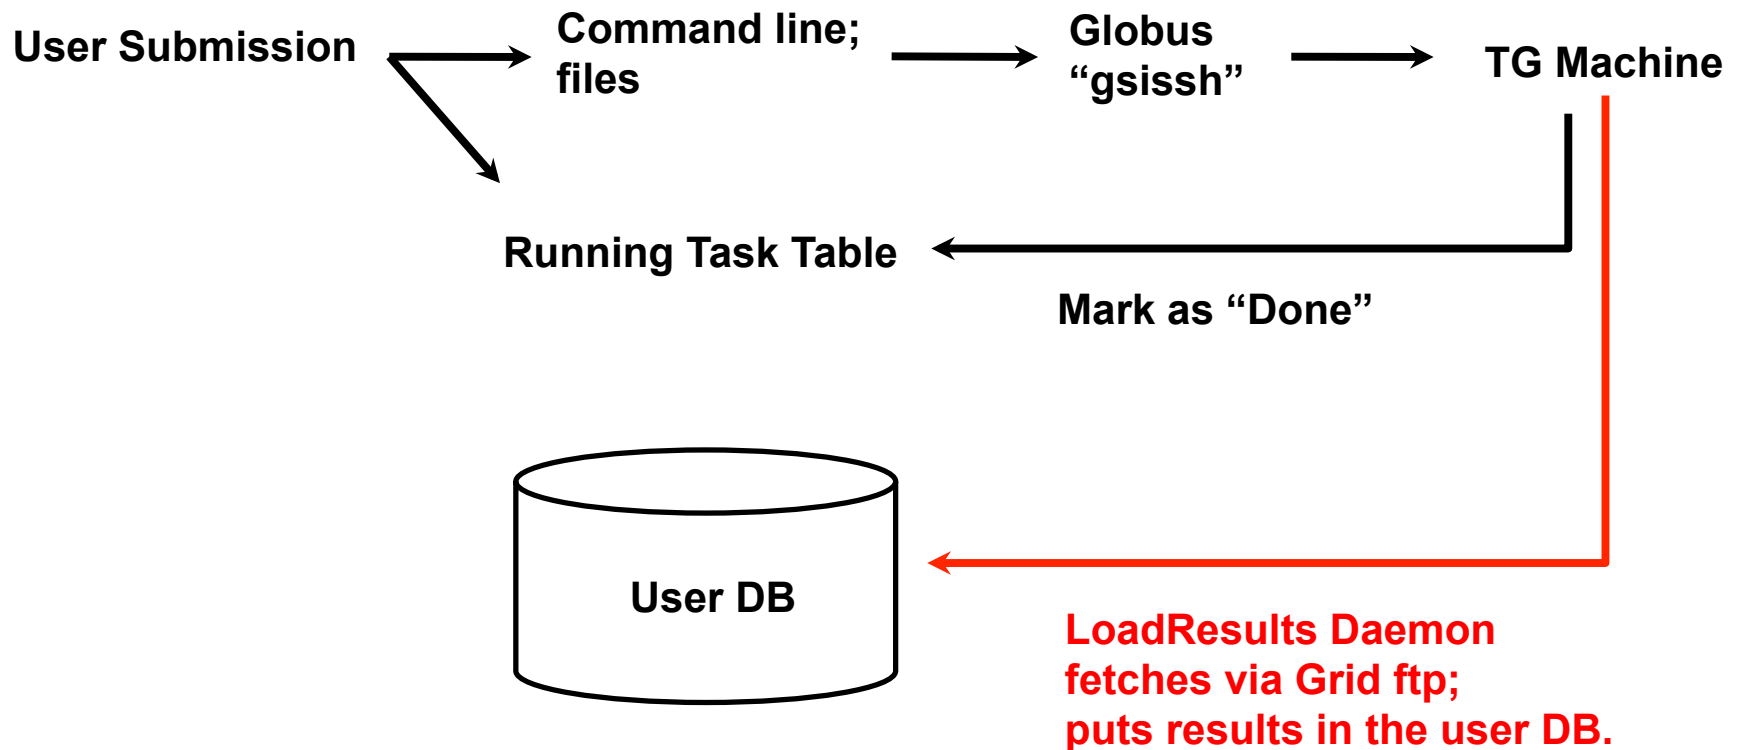


Normal Operation



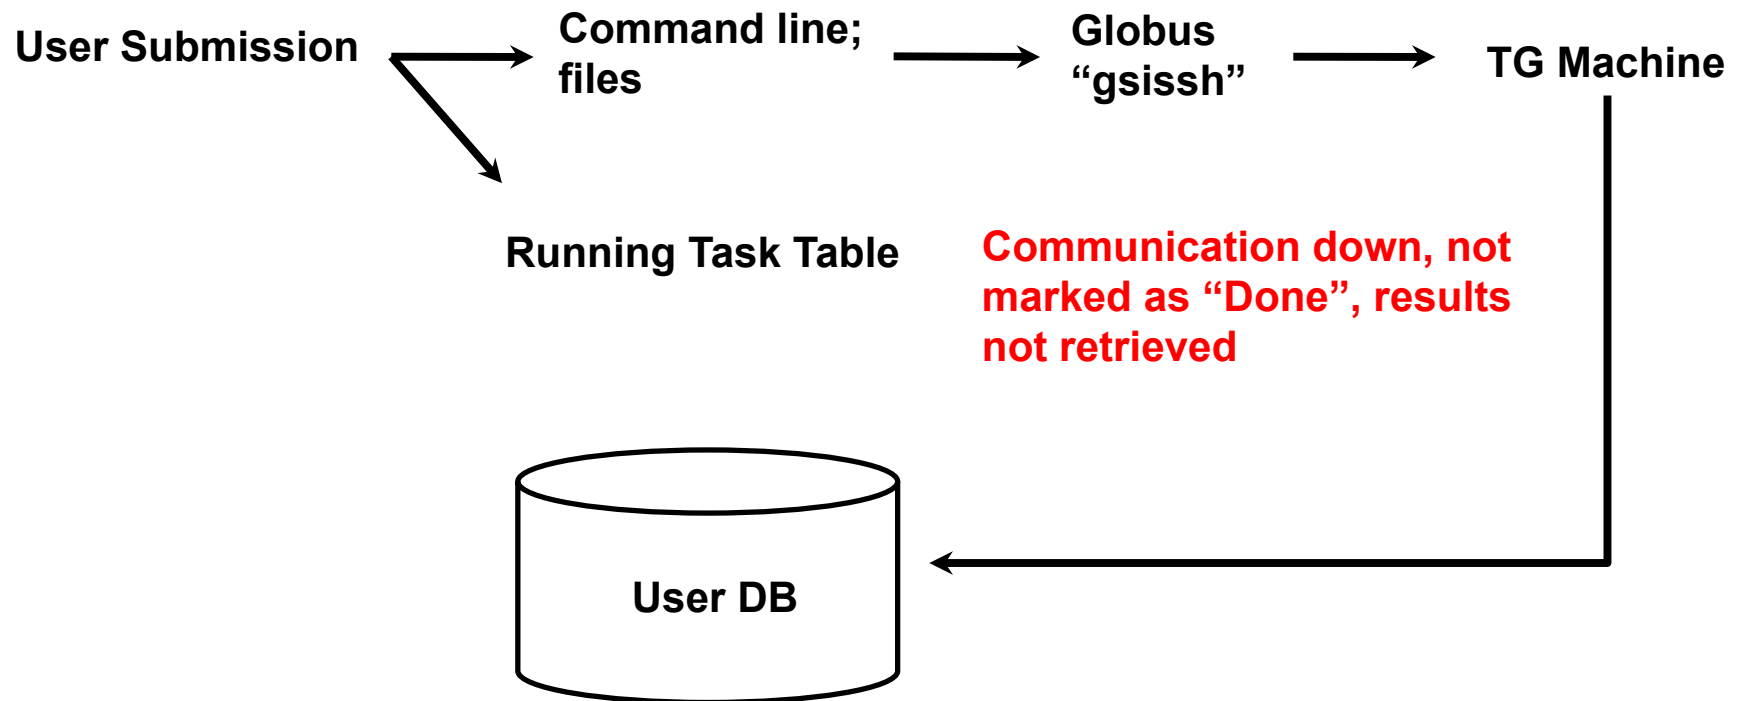


Normal Operation



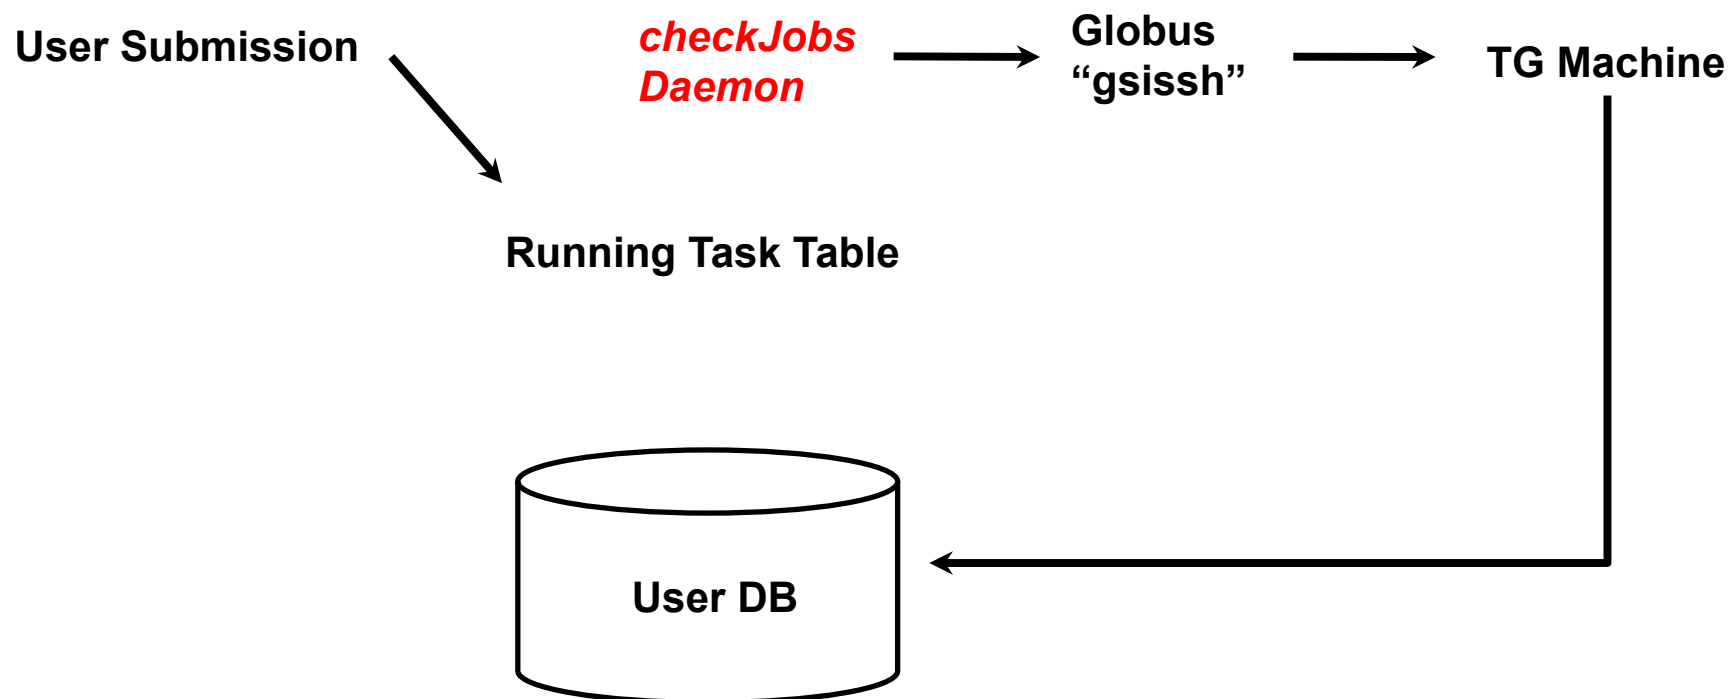


Abnormal Operation



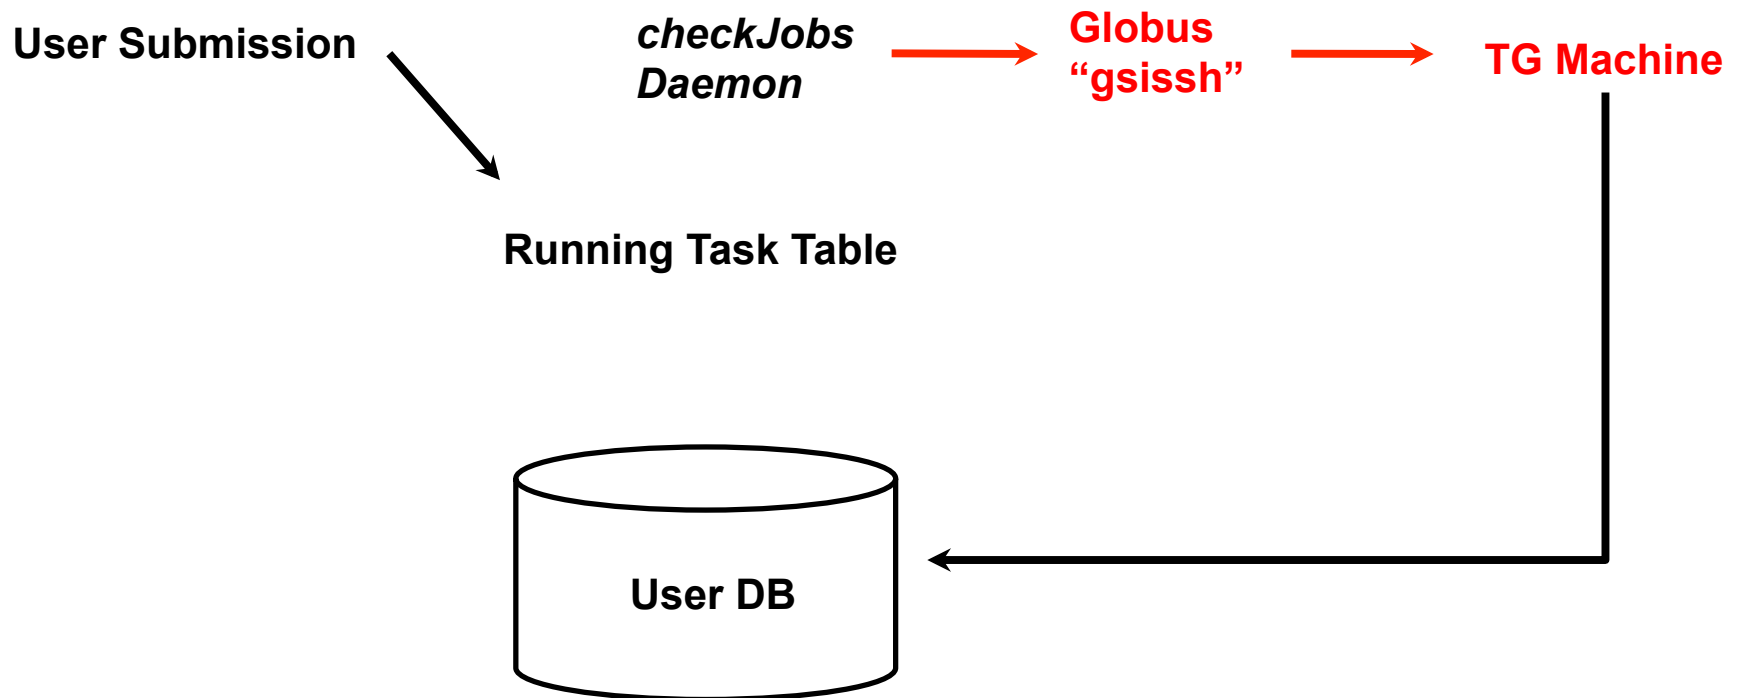


Abnormal Operation



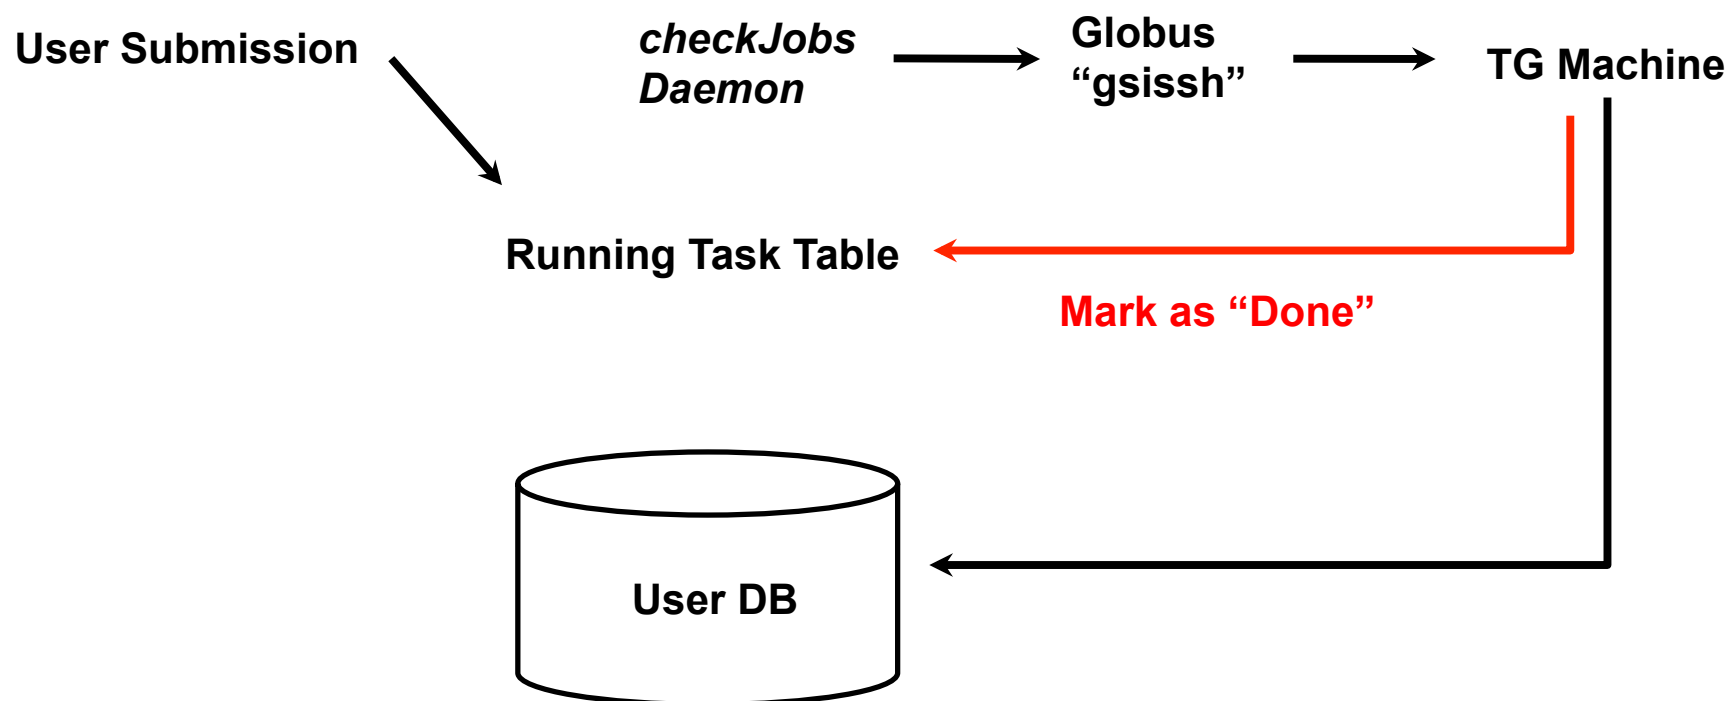


Abnormal Operation



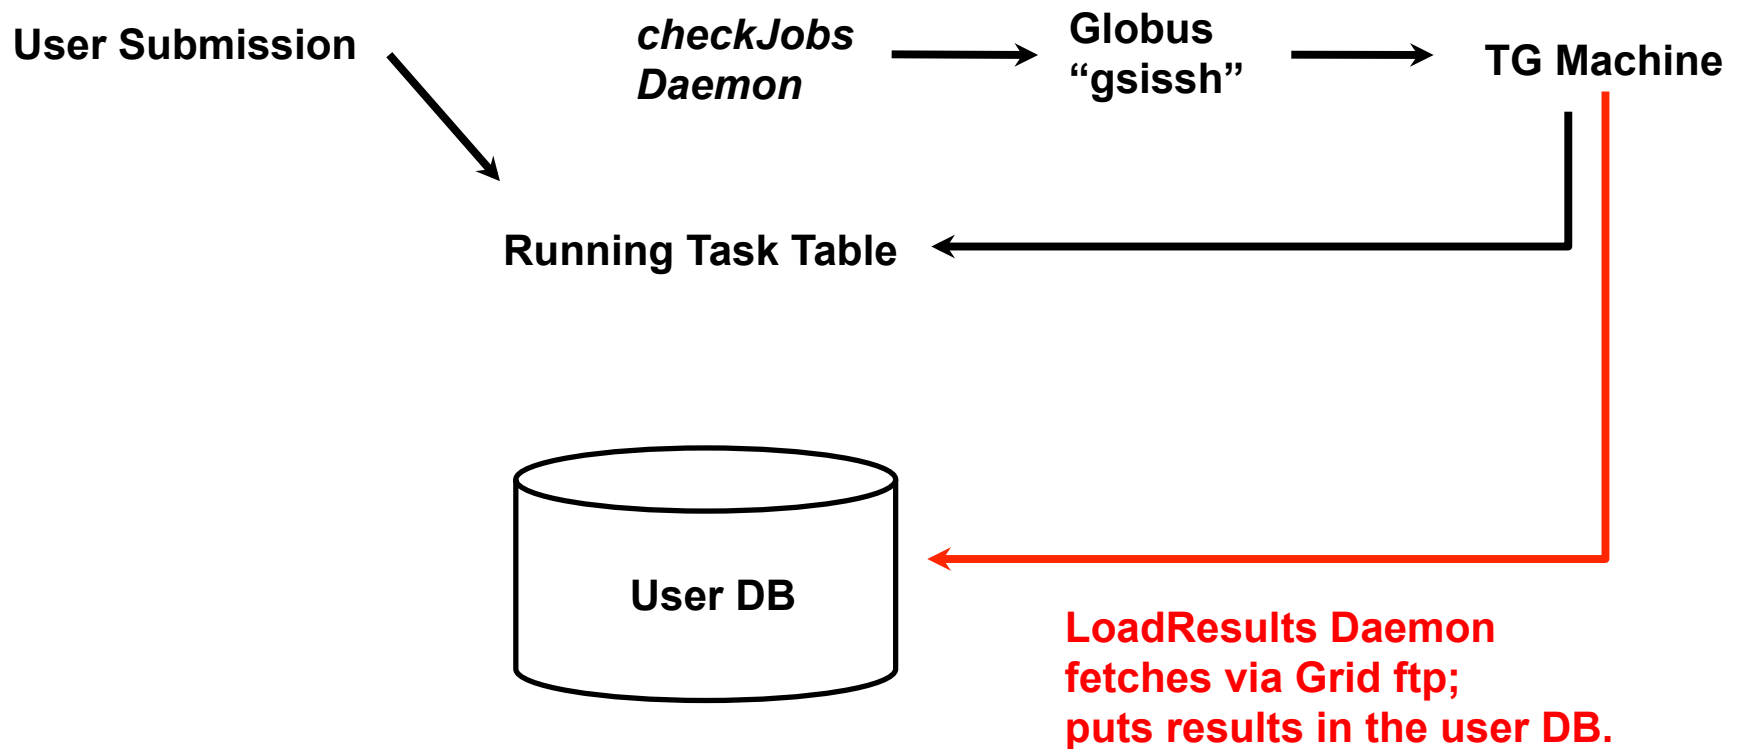


Abnormal Operation





Abnormal Operation





CIPRES

Cyberinfrastructure for
Phylogenetic Research



Jobs Saved by the GSISSH / Task Table System

	<u>SEPT</u>	<u>OCT</u>
Jobs Saved	159*	266*

*** 7% of all submitted jobs**



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



What happens if you build it and too many people come???

- **make sure resource use is efficient**
- **make sure resource use delivers impact**
- **implement policies to maximize access for all**



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



Make sure resource use delivers impact: Usage and Outcomes



POWERED BY  TeraGrid™

SDSC




CIPRES

Cyberinfrastructure for
Phylogenetic Research



CIPRES Science Gateway Usage Dec 2009 – April 2011



POWERED BY  TeraGrid

SDSC




CIPRES

Cyberinfrastructure for
Phylogenetic Research



CIPRES Science Gateway Usage Dec 2009 – April 2011



**4X increase in
users/month**



POWERED BY  TeraGrid

SDSC

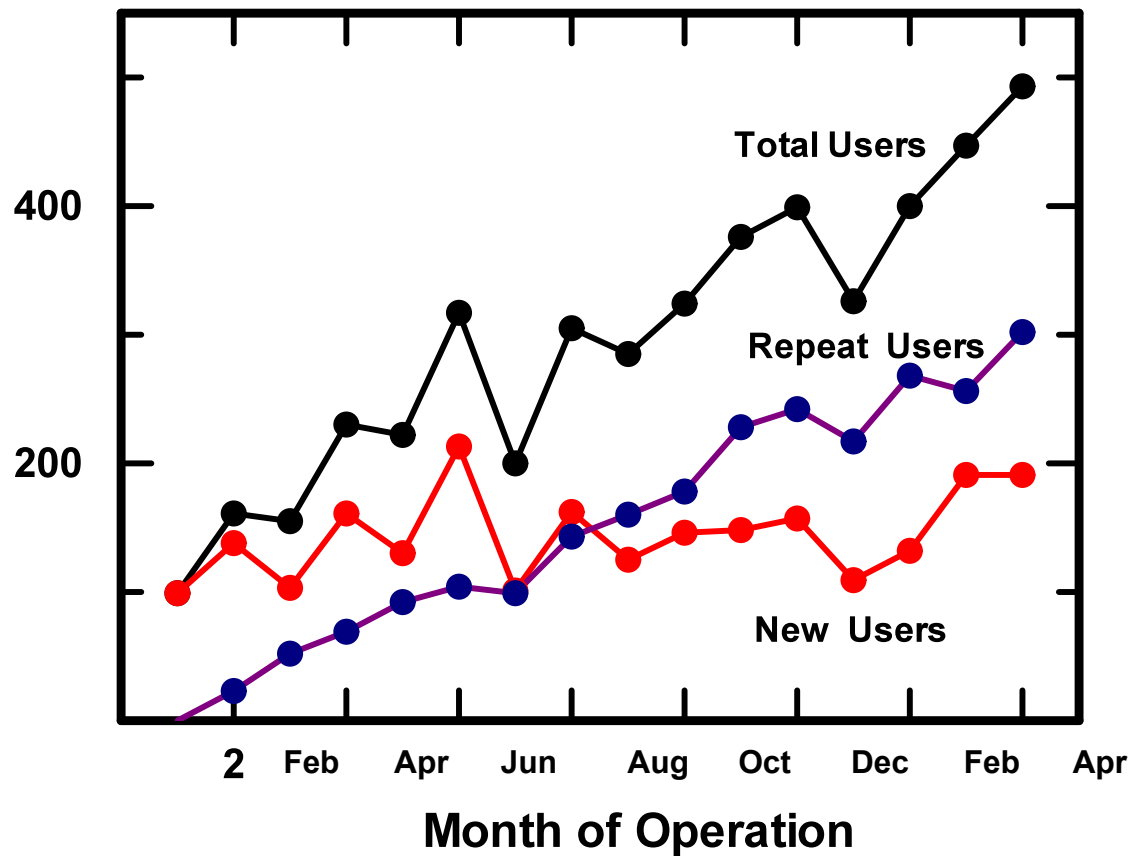


CIPRES

Cyberinfrastructure for
Phylogenetic Research



CIPRES Science Gateway Usage Dec 2009 – April 2011



**4X increase in
users/month**

**2456 new TG
users**



POWERED BY  TeraGrid

SDSC




CIPRES

Cyberinfrastructure for
Phylogenetic Research



CIPRES Science Gateway Usage Dec 2009 – April 2011



**4X increase in
users/month**

**2456 new TG
users**

**At least 100
new TG users
in each month**



POWERED BY  TeraGrid

SDSC




CIPRES

Cyberinfrastructure for
Phylogenetic Research



CIPRES Science Gateway Usage Dec 2009 – April 2011




POWERED BY  TeraGrid™

SDSC



CIPRES Science Gateway Usage Dec 2009 – April 2011




**4-fold increase in
jobs submitted
/month**





CIPRES Science Gateway Usage Dec 2009 – April 2011




**4-fold increase in
jobs submitted
/month**

**10+ fold increase
in SUs/month**





CIPRES Science Gateway Usage Dec 2009 – April 2011



**4-fold increase in
jobs submitted
/month**

**10+ fold increase
in SUs/month**

**45,500 total job
submissions in
17 months**





CIPRES

Cyberinfrastructure for
Phylogenetic Research



Intellectual Merit:

Publications enabled by the CIPRES Science Gateway/CIPRES Portal:

<u>Year</u>	<u>Number</u>
2011*	38
2010	74
2009	56
2008	4

*As of April 1, 2011

Publications in the pipeline:

<u>Status</u>	<u>Number</u>
In preparation	103
In review	24



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



Intellectual Merit:

“...quick and reliable results often help with brainstorming and pushing ‘things into the right direction.’ The service is invaluable.”



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



Broad Impact:

- **Used for curriculum delivery by at least 57 instructors.**
- **Jobs run for researchers in 23/29 EPSCOR states.**
- **Routine submissions from Harvard, Berkeley, Stanford.....**
- **76% of users are in the US or have a collaborator in the US**
- **In Q1, 2011, 28% of all TG users who ran jobs ran them from the CSG**



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



Broad Impact:

“The wall time for these data sets are 30 min (2000 BS reps) [on the CSG], but locally it would take 24 hours for each run. So, instead of 10 days to complete all the analyses with a few local computers, I have publication quality trees overnight.”



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



Make sure resource use delivers impact: Usage and Outcomes

We know how many jobs are *submitted*, how many of these jobs produce useful outcomes?

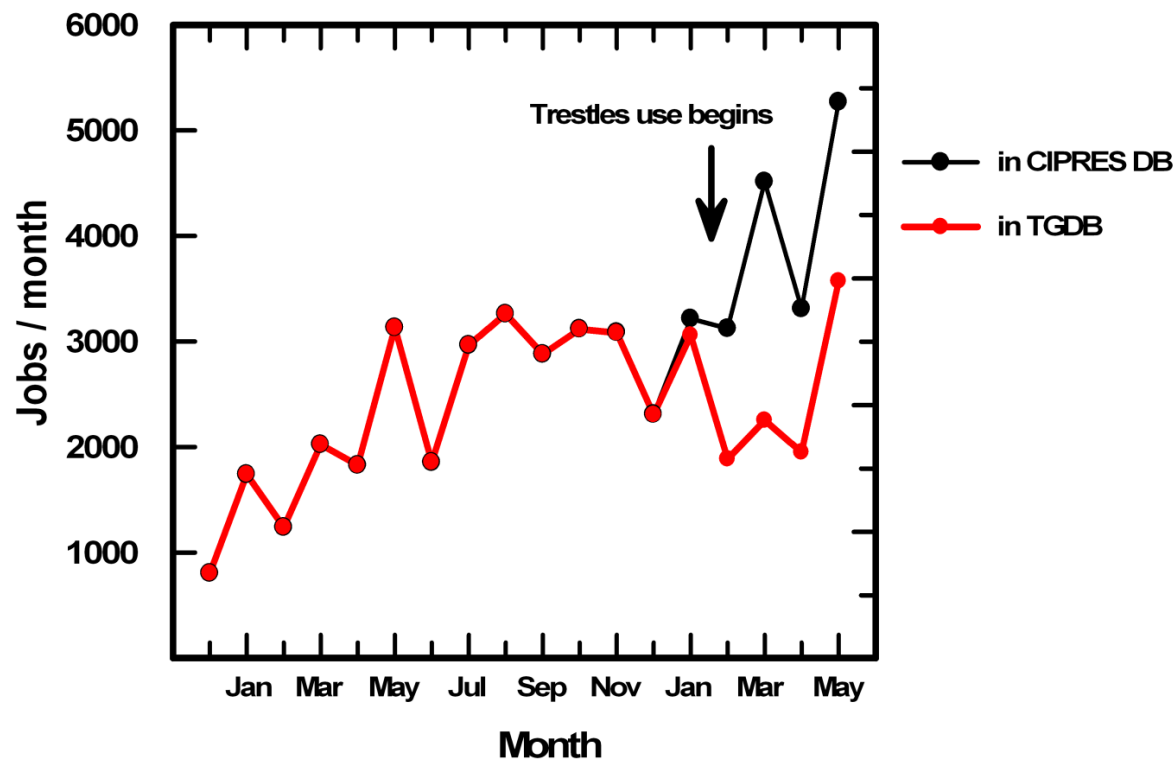


POWERED BY  TeraGrid™

SDSC

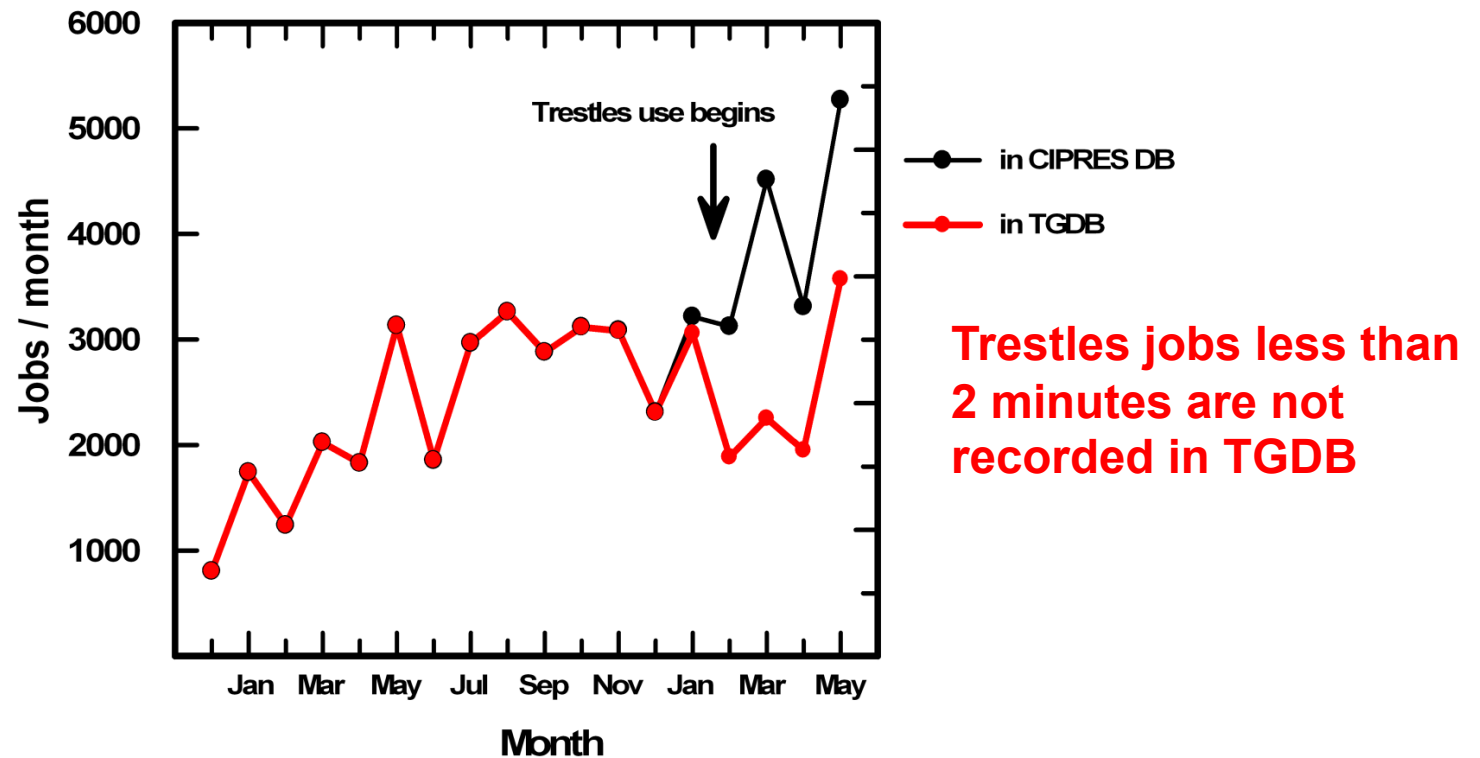


Not All Job Submissions to *Trestles* Appear in the TGDB



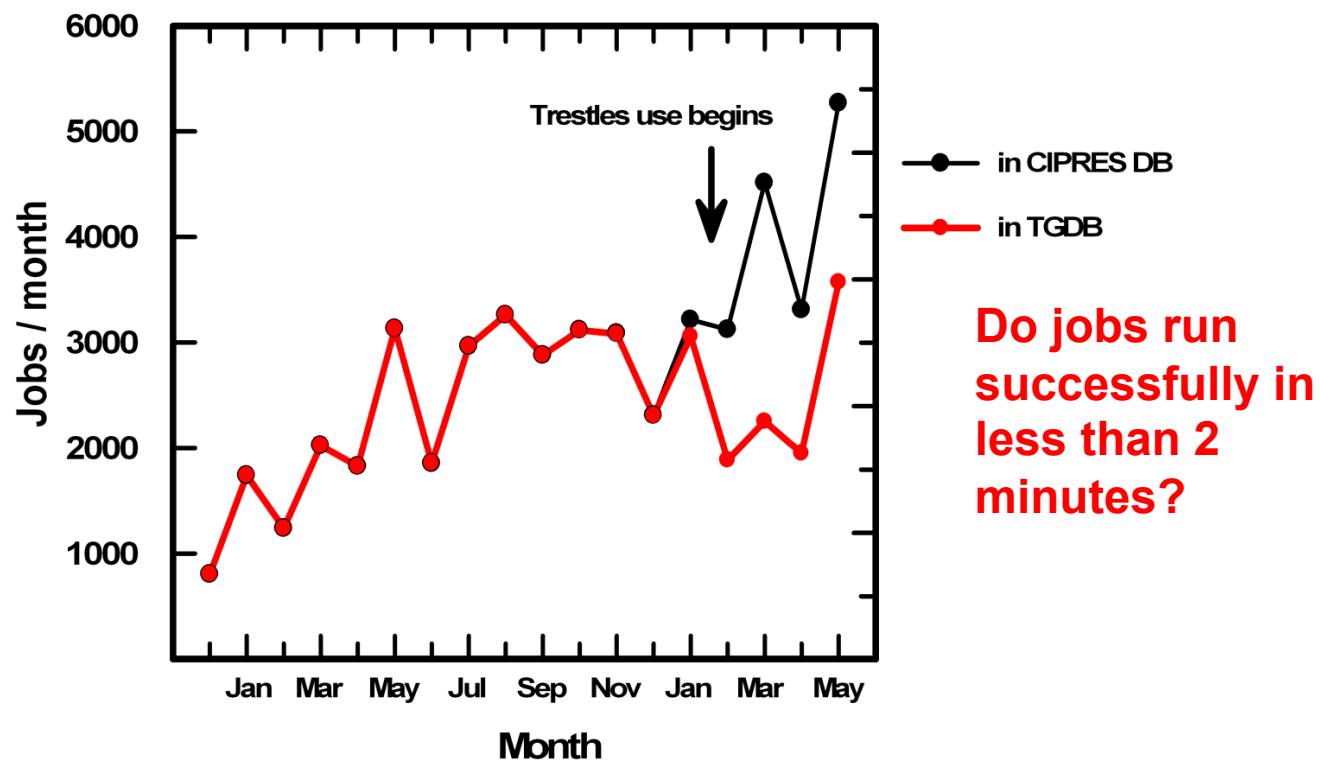


Not All Job Submissions to *Trestles* Appear in the TGDB





Not All Job Submissions to *Trestles* Appear in the TGDB





Categorization of RAxML and MrBayes jobs by outcome and duration*

	RAxML		MrBayes		
	Jobs	Fraction	Jobs	Fraction	
Failed		885	0.394	257	0.230
Succeeded		1363	0.606	860	0.770
< 2 minutes	521	0.232	18	0.016	
≥ 2 minutes	842	0.375	842	0.754	

* Jobs run between 3/16/2011 and 4/17/2011.





Categorization of RAxML and MrBayes jobs by outcome and duration*

	RAxML		MrBayes	
	Jobs	Fraction	Jobs	Fraction
Failed		885		0.394
			257	0.230
Succeeded		1363		0.606
			860	0.770
< 2 minutes	521	0.232	18	0.016
≥ 2 minutes	842	0.375	842	0.754

* Jobs run between 3/16/2011 and 4/17/2011.





Categorization of RAxML and MrBayes jobs by outcome and duration*

	RAxML		MrBayes		
	Jobs	Fraction	Jobs	Fraction	
Failed		885	0.394	257	0.230
Succeeded		1363	0.606	860	0.770
< 2 minutes	521	0.232	18	0.016	
≥ 2 minutes	842	0.375	842	0.754	

* Jobs run between 3/16/2011 and 4/17/2011.





Categorization of CSG users by job outcome and duration *

	Users	Fraction
Only Failed	50	0.13
Run Succeeded	345	0.87
Only < 2 minutes	61	0.15
≥ 2 minutes**	284	0.72

*Jobs between 3/17/2011 and 4/17/2011

** The user ran at least one job that produced results and recorded an SU charge.





Categorization of CSG users by job outcome and duration *

	Users	Fraction	
Only Failed	50	0.13	
Run Succeeded	345	0.87	
Only < 2 minutes	61	0.15	15% of all users only submitted short jobs
≥ 2 minutes**	284	0.72	

*Jobs between 3/17/2011 and 4/17/2011

** The user ran at least one job that produced results and recorded an SU charge.





Categorization of CSG users by job outcome and duration *

	Users	Fraction	
Only Failed	50	0.13	13% of all users did not submit a successful job
Run Succeeded	345	0.87	
Only < 2 minutes	61	0.15	
≥ 2 minutes**	284	0.72	

*Jobs between 3/17/2011 and 4/17/2011

** The user ran at least one job that produced results and recorded an SU charge.





CIPRES

Cyberinfrastructure for
Phylogenetic Research



What happens if you build it and too many people come???

- **make sure resource use is efficient**
- **make sure resource use delivers impact**
- **implement policies to maximize access for all**




POWERED BY  TeraGrid™

SDSC



How much growth can we manage?





CIPRES

Cyberinfrastructure for
Phylogenetic Research



Establish a Fair Use Policy

- How many SUs should each user be permitted?
- How much influence should affiliation with a US institution have on resource allocation?




POWERED BY  TeraGrid™

SDSC

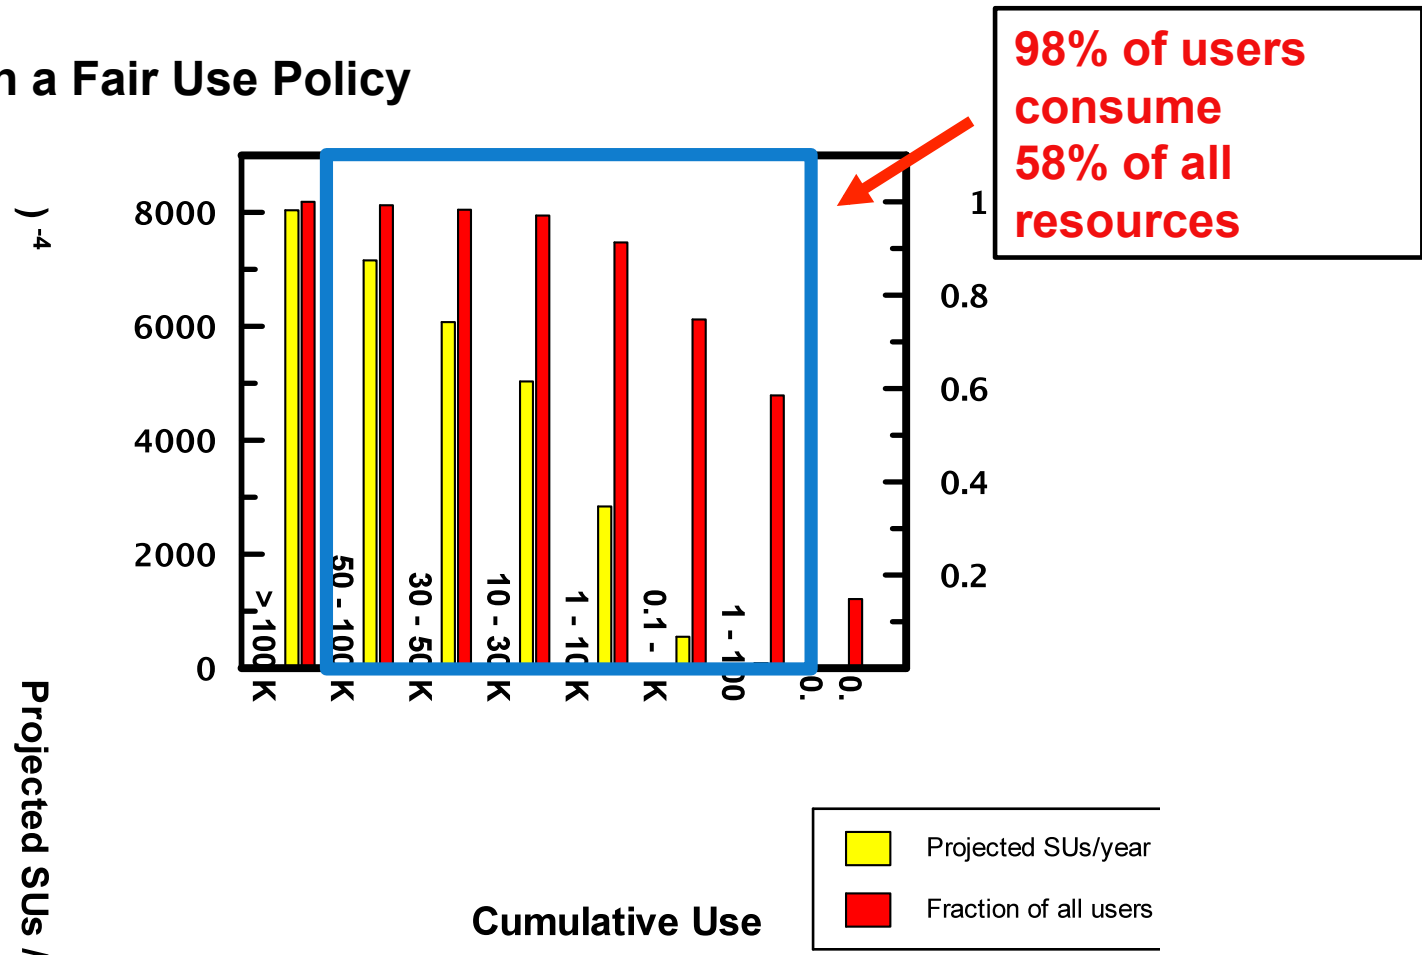


Establish a Fair Use Policy






Establish a Fair Use Policy



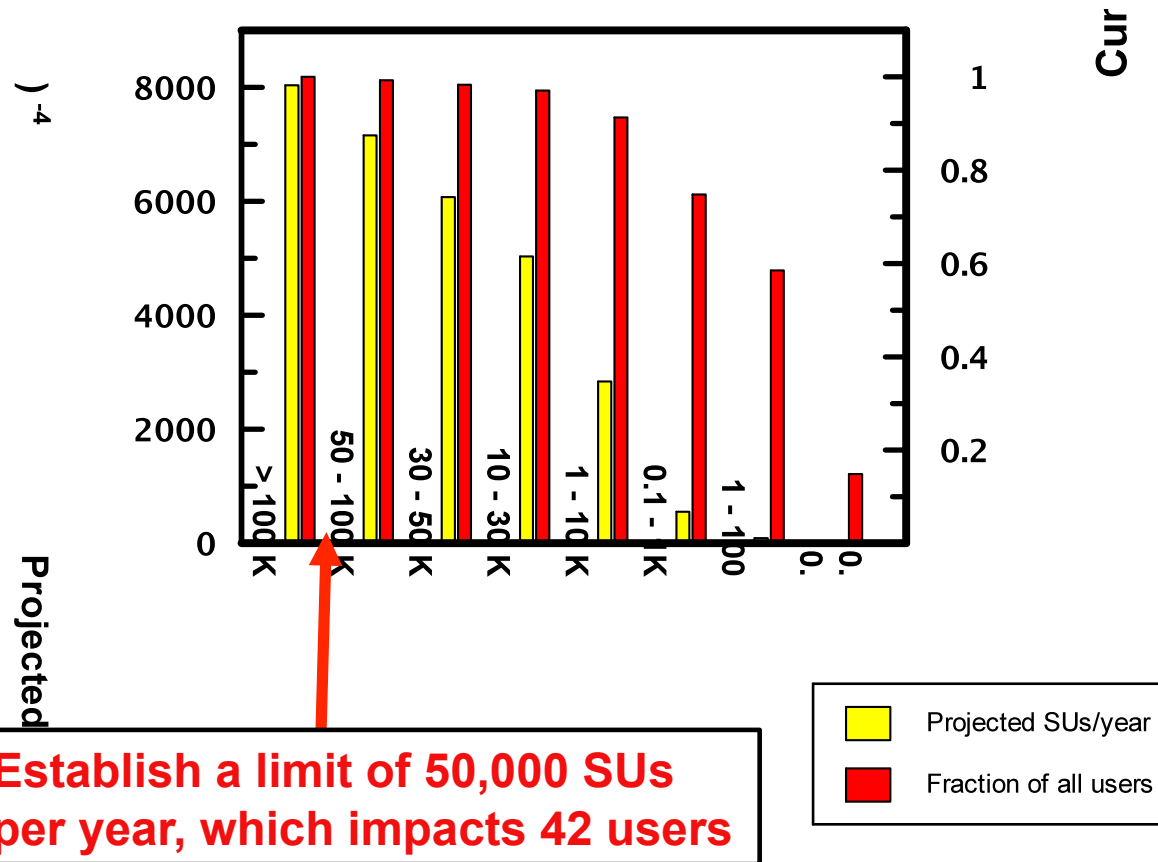


Establish a Fair Use Policy





Establish a Fair Use Policy





CIPRES

Cyberinfrastructure for
Phylogenetic Research



Establish a Fair Use Policy

- Users are permitted to use 50,000 SUs annually from the community allocation. Anyone from any institution can register.
- Users who wish to use in excess of this amount can apply for a personal allocation. **This would include users at US institutions and those who have collaborators at a US institution.** User can access their allocation using the CSG.



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



Tools required to implement the CIPRES SG Fair Use Policy:

- **ability to monitor usage by each account automatically (complete)**
- **ability to halt submissions from a given user account (complete)**
- **ability to charge to a user's personal allocation (complete)**
- **ability for users to track their SU consumption (complete)**
- **ability to forecast SU cost of a job for users (in progress)**



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



Next Steps:

- **Deploy the new user management tools in production.**
- **Implement the new Fair Usage Policy.**
- **Add new codes: BEAST, RAxML Light, BEST.**
- **Partner with the iPlant Project as a service provider.**



POWERED BY  TeraGrid™

SDSC

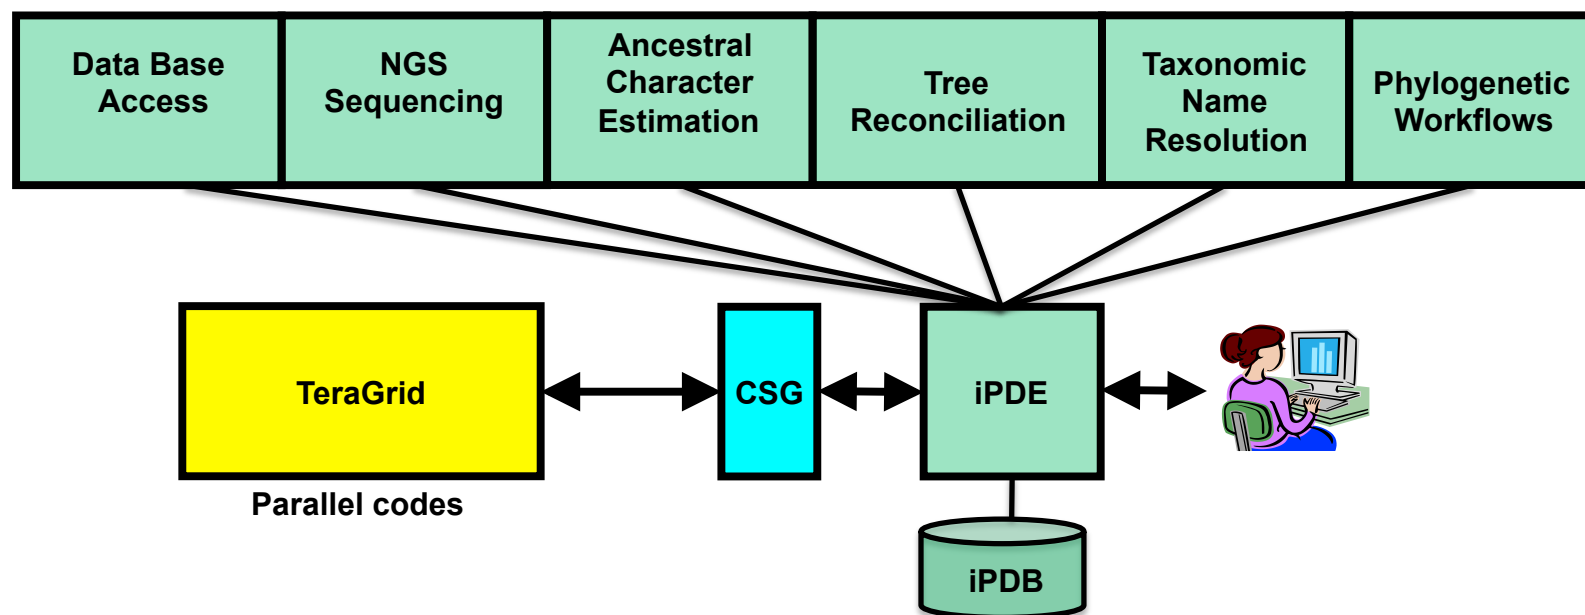


CIPRES

Cyberinfrastructure for
Phylogenetic Research



Next Steps:



POWERED BY  TeraGrid™

SDSC



CIPRES

Cyberinfrastructure for
Phylogenetic Research



Acknowledgements:

CIPRES Science Gateway

Terri Schwartz

TeraGrid Hybrid Code Development

**Wayne Pfeiffer
Alexandros Stamatakis**

TeraGrid Implementation Support

**Nancy Wilkins-Diehr
Doru Marcusiu
Leo Carson**

Workbench Framework:

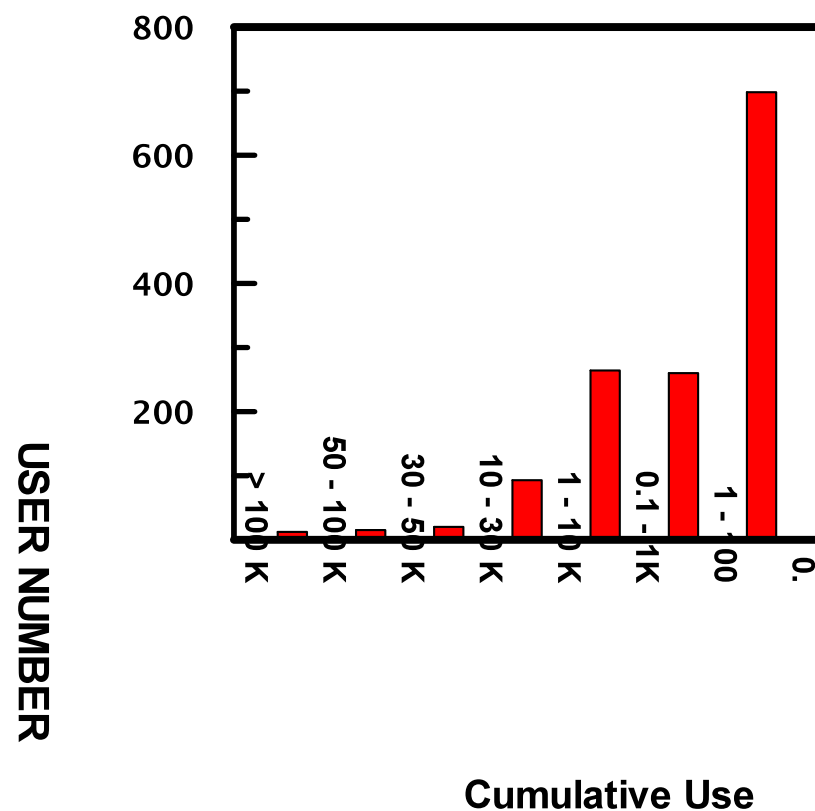
**Paul Hoover
Lucie Chan**



SDSC



Make sure resource use delivers impact: analyze resource consumption




* Reporting Period: Sept, 2010 – May, 2011





Make sure resource use delivers impact: analyze resource consumption

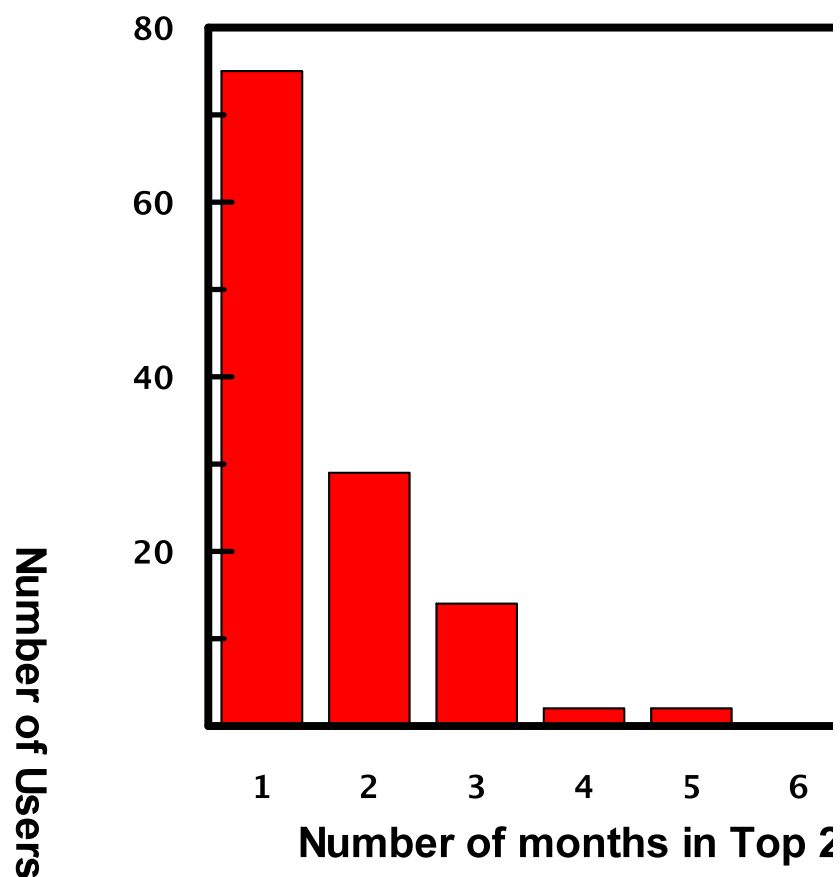


* Reporting Period: Sept, 2010 – May, 2011





Identify the TOP 20 Users in each month.

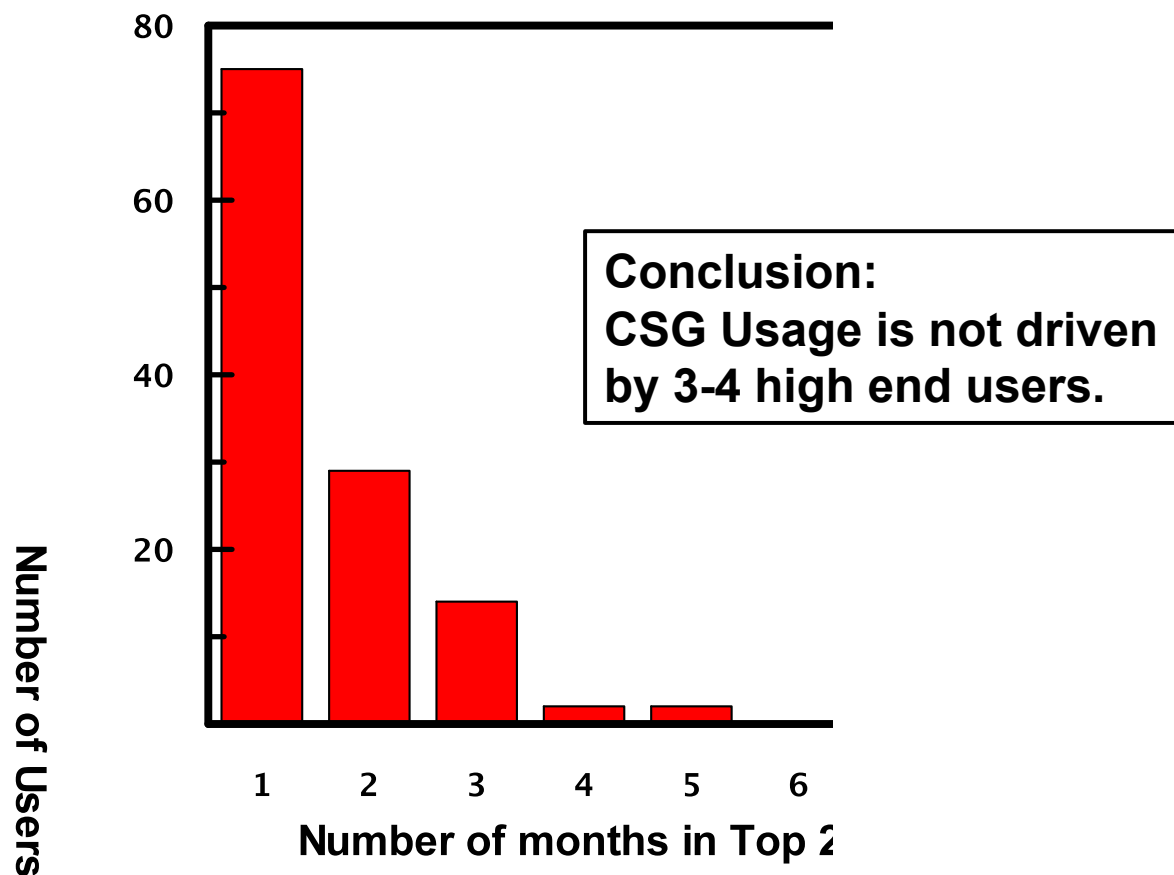


* Reporting Period: Sept, 2010 – May, 2011





Identify the TOP 20 Users in each month.



* Reporting Period: Sept, 2010 – May, 2011





Make sure resource use delivers impact: analyze resource consumption

SUs	% of Users	% total SU
0 – 30 K	97	45
30 – 300,000 K	3	55
> 300,000 K	0	0

* Reporting Period: Sept, 2010 – May, 2011





Make sure resource use delivers impact: analyze resource consumption

SUs	% of Users	% total SU
0 – 30 K	97	45
30 – 300,000 K	3	55
> 300,000 K	0	0

**50% of users in
this bin are from
institutions
outside the US**

* Reporting Period: Sept, 2010 – May, 2011





Make sure resource use delivers impact: Establish a Fair Use Policy

SUs	Number of Users	% total SU	% per user
< 100	936 (58.5%)	0.2	0.0003
100 - 999	260 (16.3%)	1.6	0.005
1 – 10 K	264 (16.5%)	16	0.06
10 – 30 K	93 (5.8%)	27	0.29
30 – 50 K	20 (1.3%)	14	0.69
50 – 100 K	15 (0.9%)	16	1.06
> 100 K	12 (0.8%)	26	2.14





Make sure resource use delivers impact: Establish a Fair Use Policy

SUs	Number of Users	% total SU	% per user
< 100	936 (58.5%)	0.2	0.0003
100 - 999	260 (16.3%)		
1 - 10 K	264 (16.5%)		
10 - 30 K	93 (5.8%)		
30 - 50 K	20 (1.3%)	14	0.69
50 - 100 K	15 (0.9%)	16	1.06
> 100 K	12 (0.8%)	26	2.14

Establish 50,000 SUs as the limit for usage from the CSG community account. (Same as Trestles startup account).





Make sure resource use delivers impact: Establish a Fair Use Policy

SUs	Number of Users	% total SU	% per user
< 100	936 (58.5%)	0.2	0.0003
100 - 999	260 (16.3%)	1.6	0.005
1 - 10 K	264 (16.5%)	16	0.06
10 - 30 K	93 (5.8%)	27	0.29
30 - 50 K	20 (1.3%)	14	0.69
50 - 100 K	15 (0.9%)	16	1.06
> 100 K	12 (0.8%)	26	2.14

Impacts 42 Users only





Make sure resource use delivers impact: resource consumption

The screenshot shows the CIPRES Science Gateway interface. On the left is a 'Folders' sidebar with various user folders. The main area is titled 'Tasks' and displays a table of active tasks. At the top right of the task area, it shows 'Current CPU Hr Usage: 10001' with a link to 'Explain this?'. Below this, a message states 'There are currently 20 tasks in this tab. (Items 1 - 20 are shown here.)'. The task table has columns for 'Select all', 'Label', 'Tool', 'Input', 'Parameters', 'Output', 'Date Created', and 'Action'.

Select all	Label	Tool	Input	Parameters	Output	Date Created	Action
<input type="checkbox"/>	ww	RAxML-HPC2 on TG	View (1)	View (26)	View (2)	6/13/11, 18:26	View Output
<input type="checkbox"/>	ww	RAxML-HPC2 on TG	View (1)	View (26)	View (2)	6/3/11, 20:42	View Output
<input type="checkbox"/>	ww	RAxML-HPC2 on TG	View (1)	View (26)	View (2)	5/24/11, 19:19	View Output
<input type="checkbox"/>	ww	RAxML-HPC2 on TG	View (1)	View (26)	View (2)	5/18/11, 09:12	View Output

Notify users of their usage level

