



# CIPRES

Cyberinfrastructure for  
Phylogenetic Research



## A Brief Overview of CIPRES Job Management

**Mark A. Miller**

**San Diego Supercomputer Center**




**XSEDE**  
Extreme Science and Engineering  
Discovery Environment

**SDSC**



**CIPRES is a classic browser-based gateway designed to submit jobs to multiple hosts. (REST services are now available also)**





## CIPRES Use Case:

- **25,000 submissions per month (34 per hour)**
- **Some require long run times (up to 168 h max)**
- **Many codes have no restart ability (some do)**
- **Relatively few cores (2 - 64, generally)**
- **Small input data sets (1 - 50 MB, generally)**
- **Many community codes (more than we can implement)**
- **Relatively unsophisticated (from HPC point of view) users**
- **Less than 2 fte of support staff**



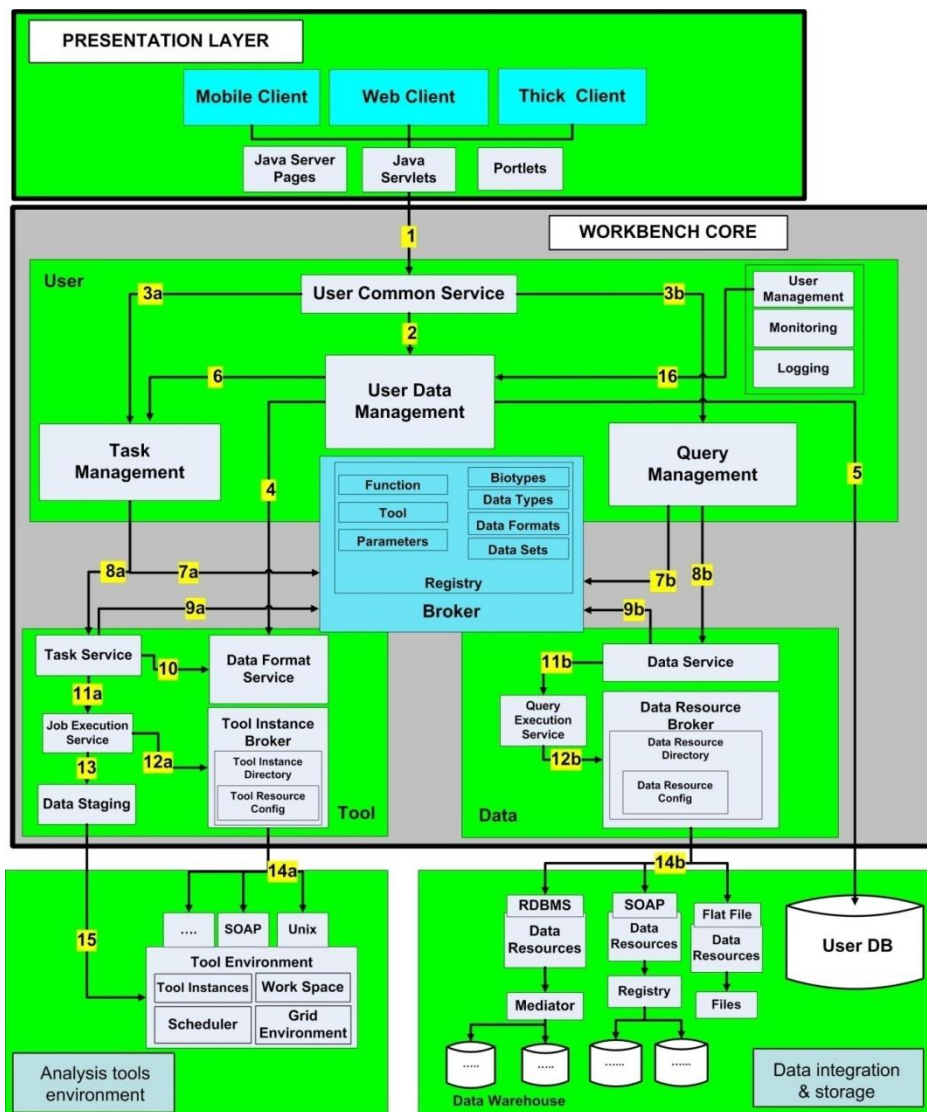


## CIPRES Use Case:

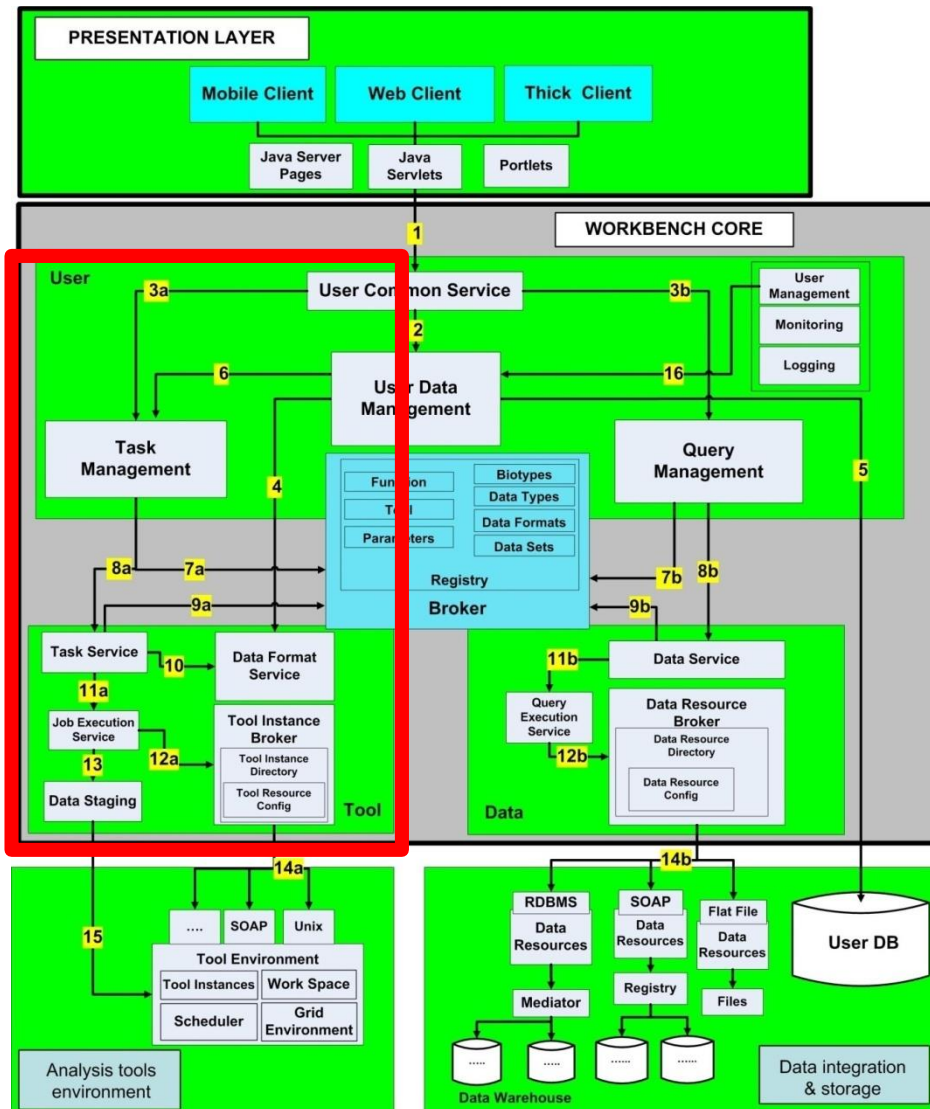
- **25,000 submissions per month (34 per hour)**
- **Some require long run times (up to 168 h max)**
- **Many codes have no restart ability (some do)**
- **Relatively few cores (2 - 64, generally)**
- **Small input data sets (1 - 50 MB, generally)**
- **Many community codes (more than we can implement)**
- **Relatively unsophisticated (from HPC point of view) users**
- **Less than 2 fte of support staff**

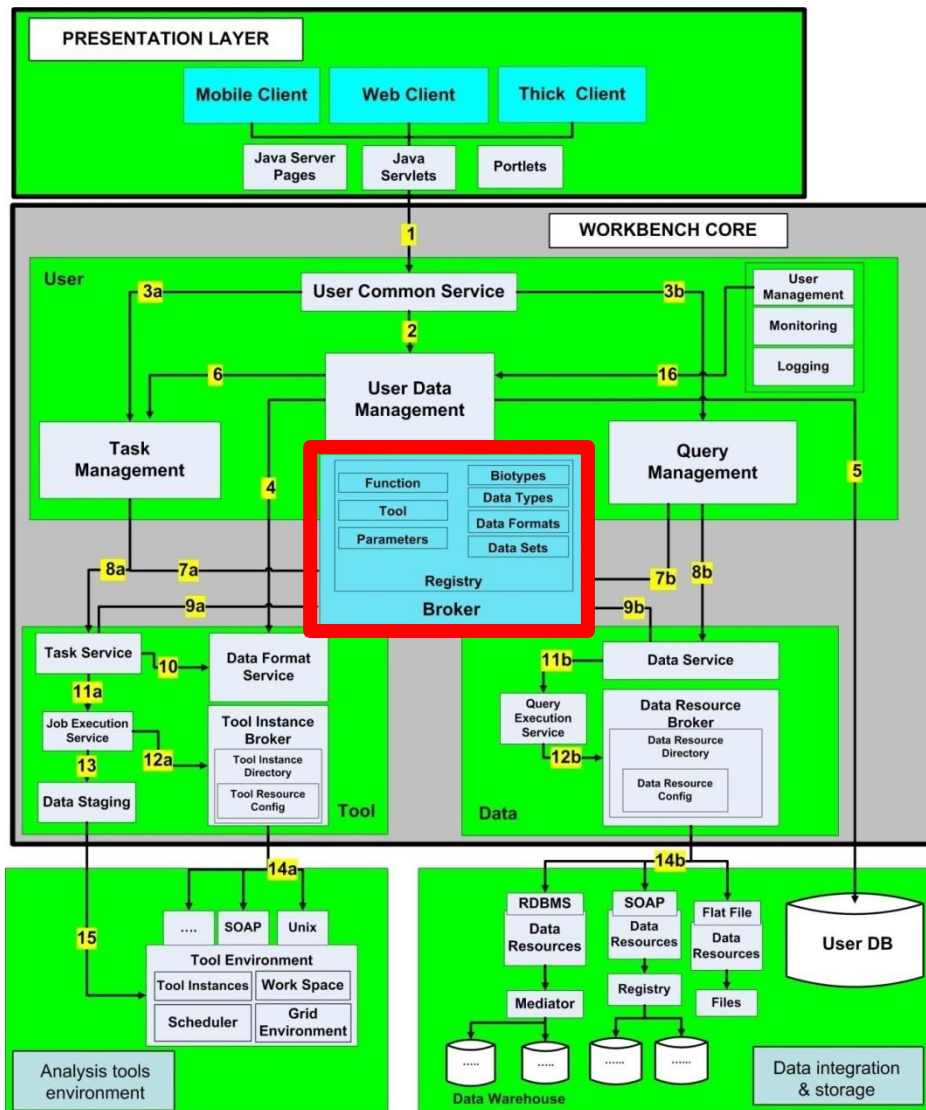
**Your use case will be different...each use case has its own issues**





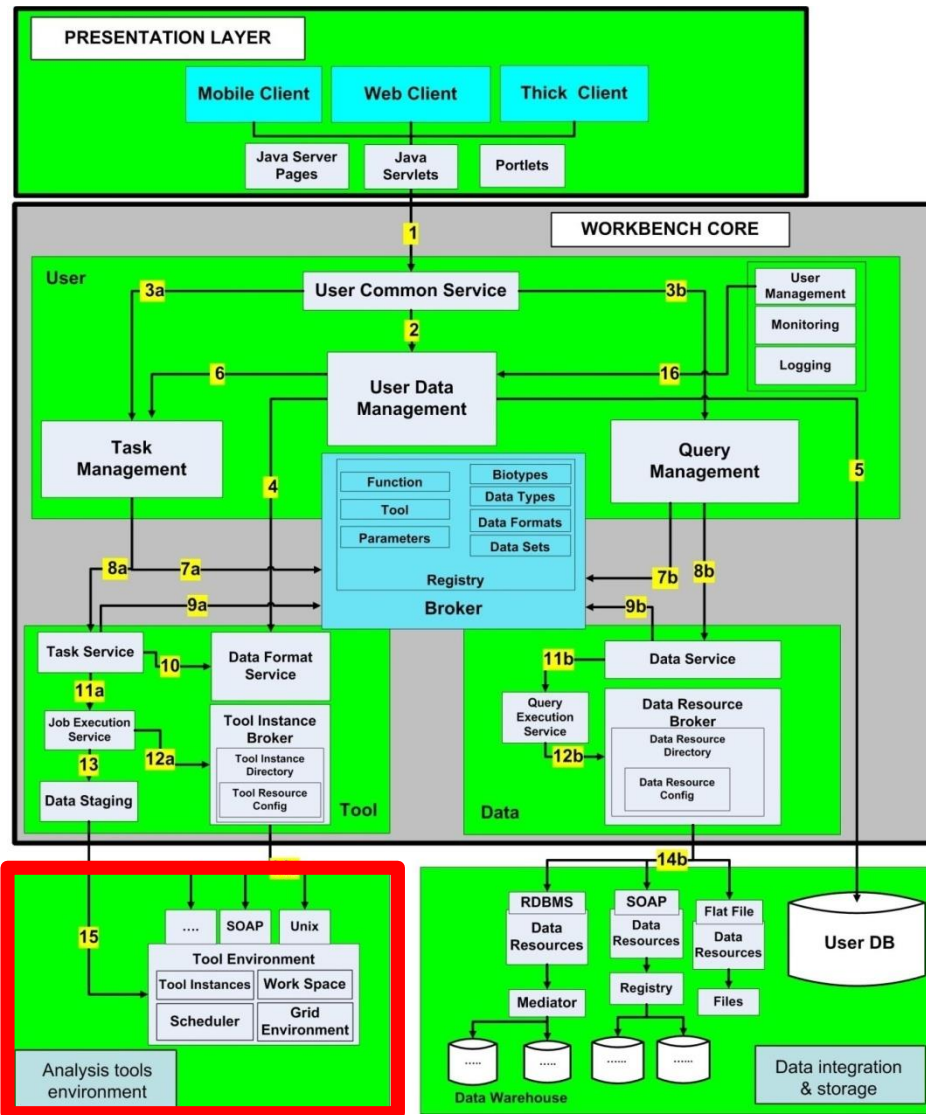
The Workbench Framework (Java) deploys generic “tasks”....





Specific information  
is coded in a  
Central Registry

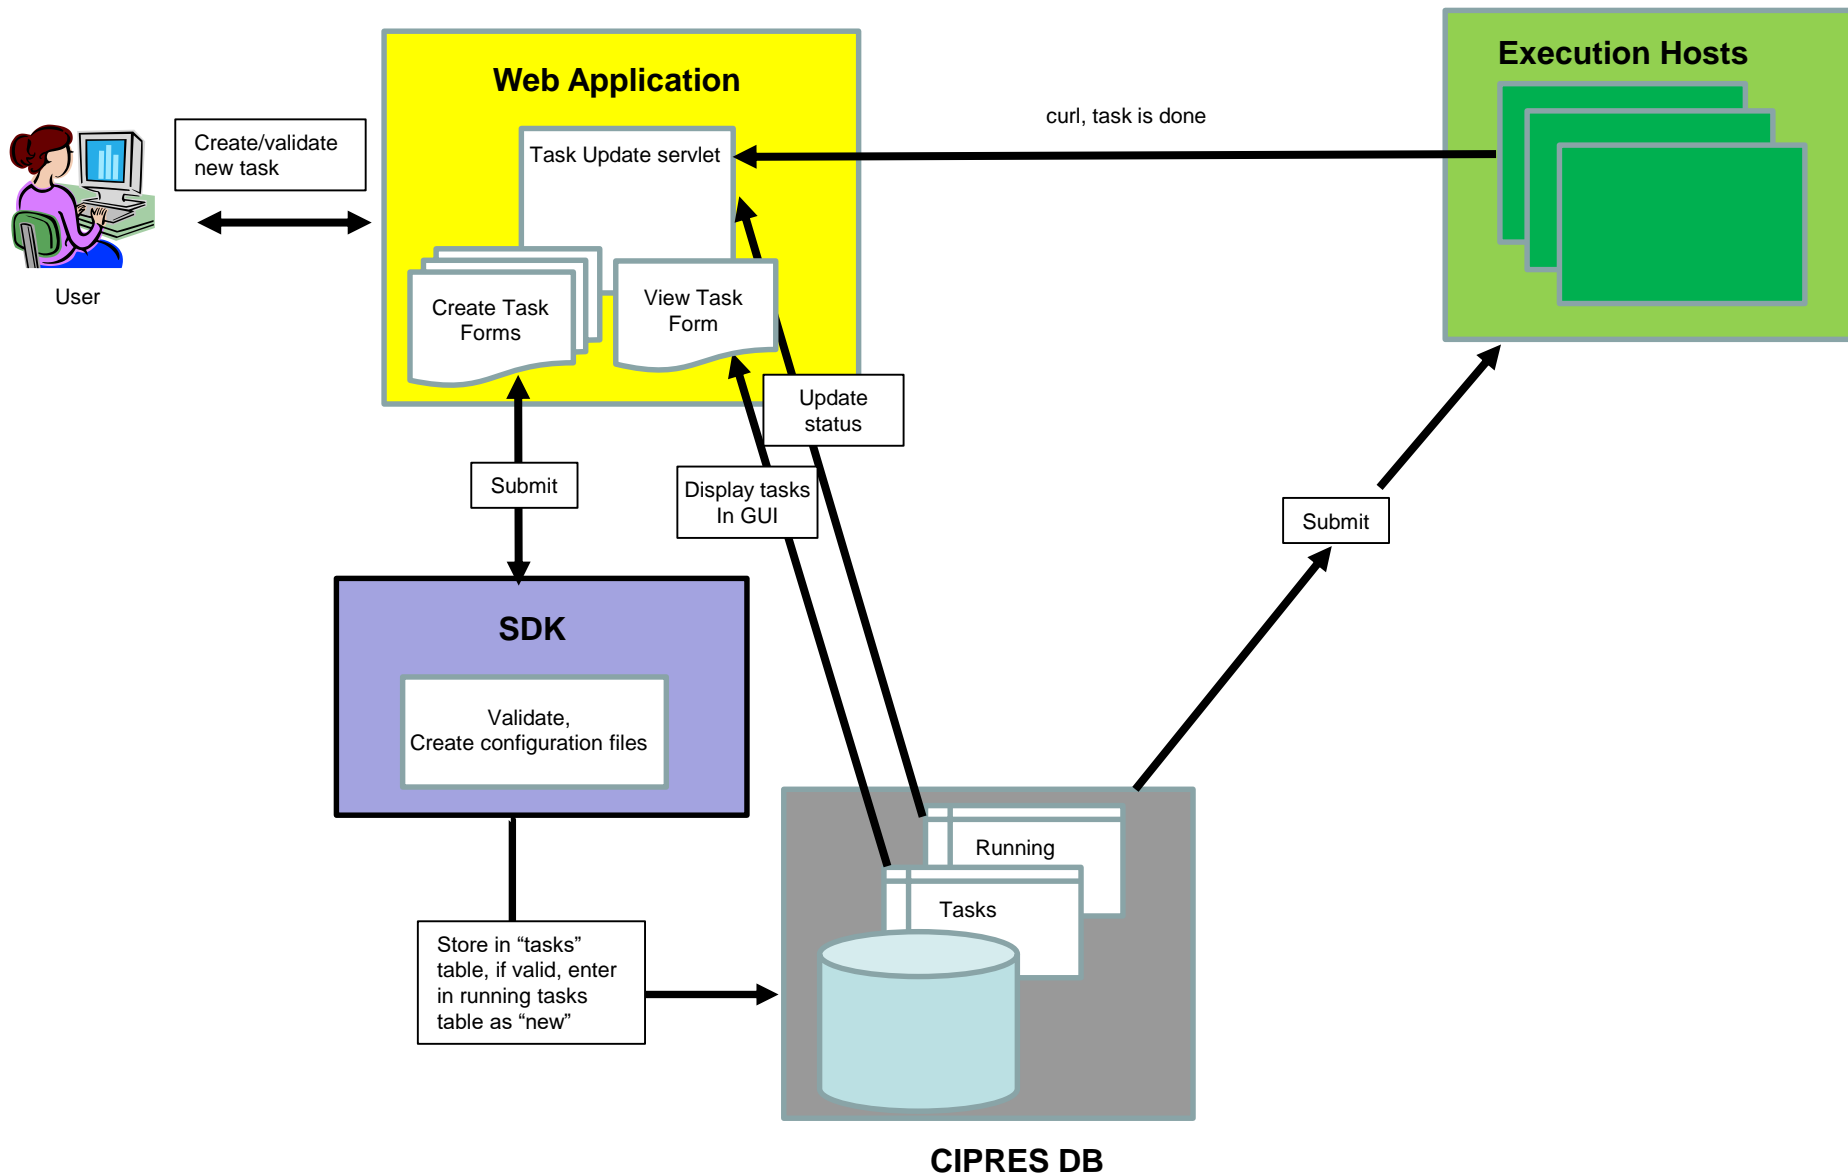


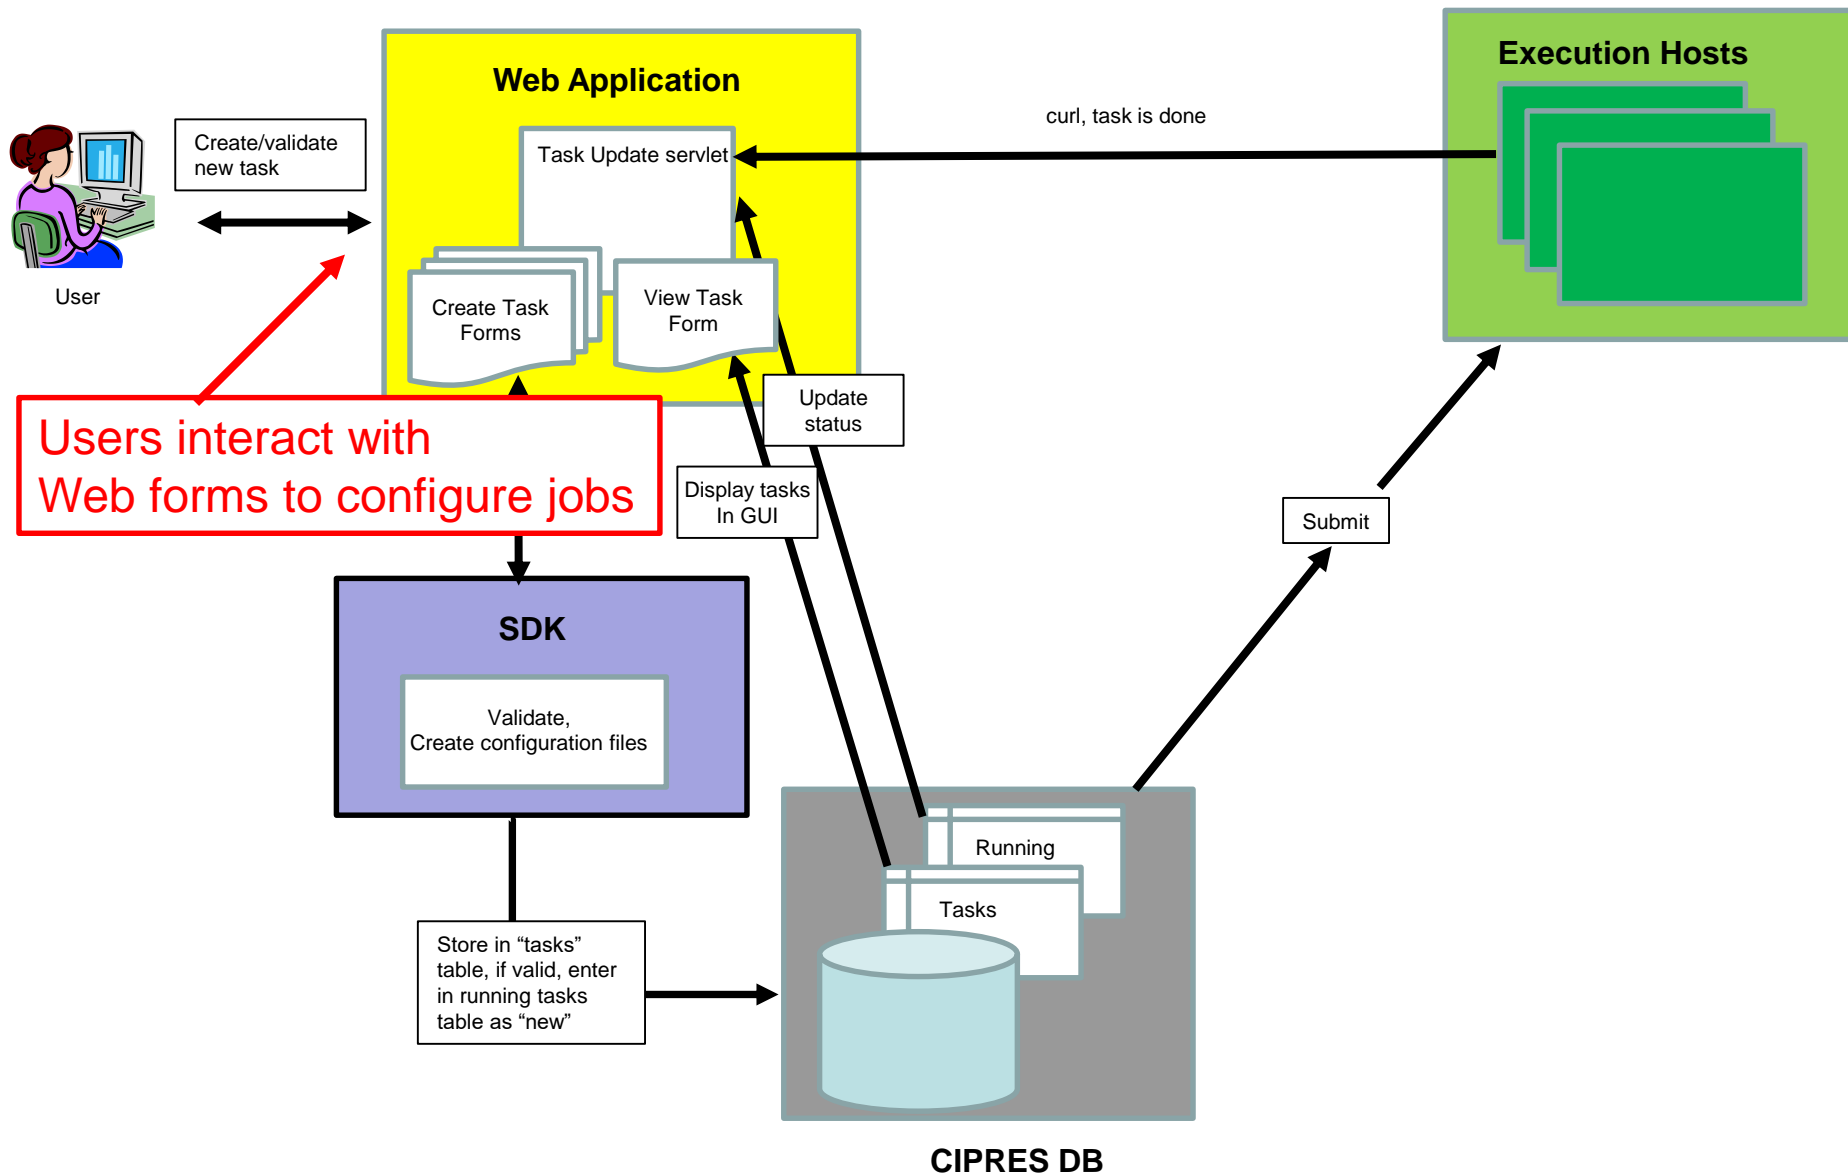


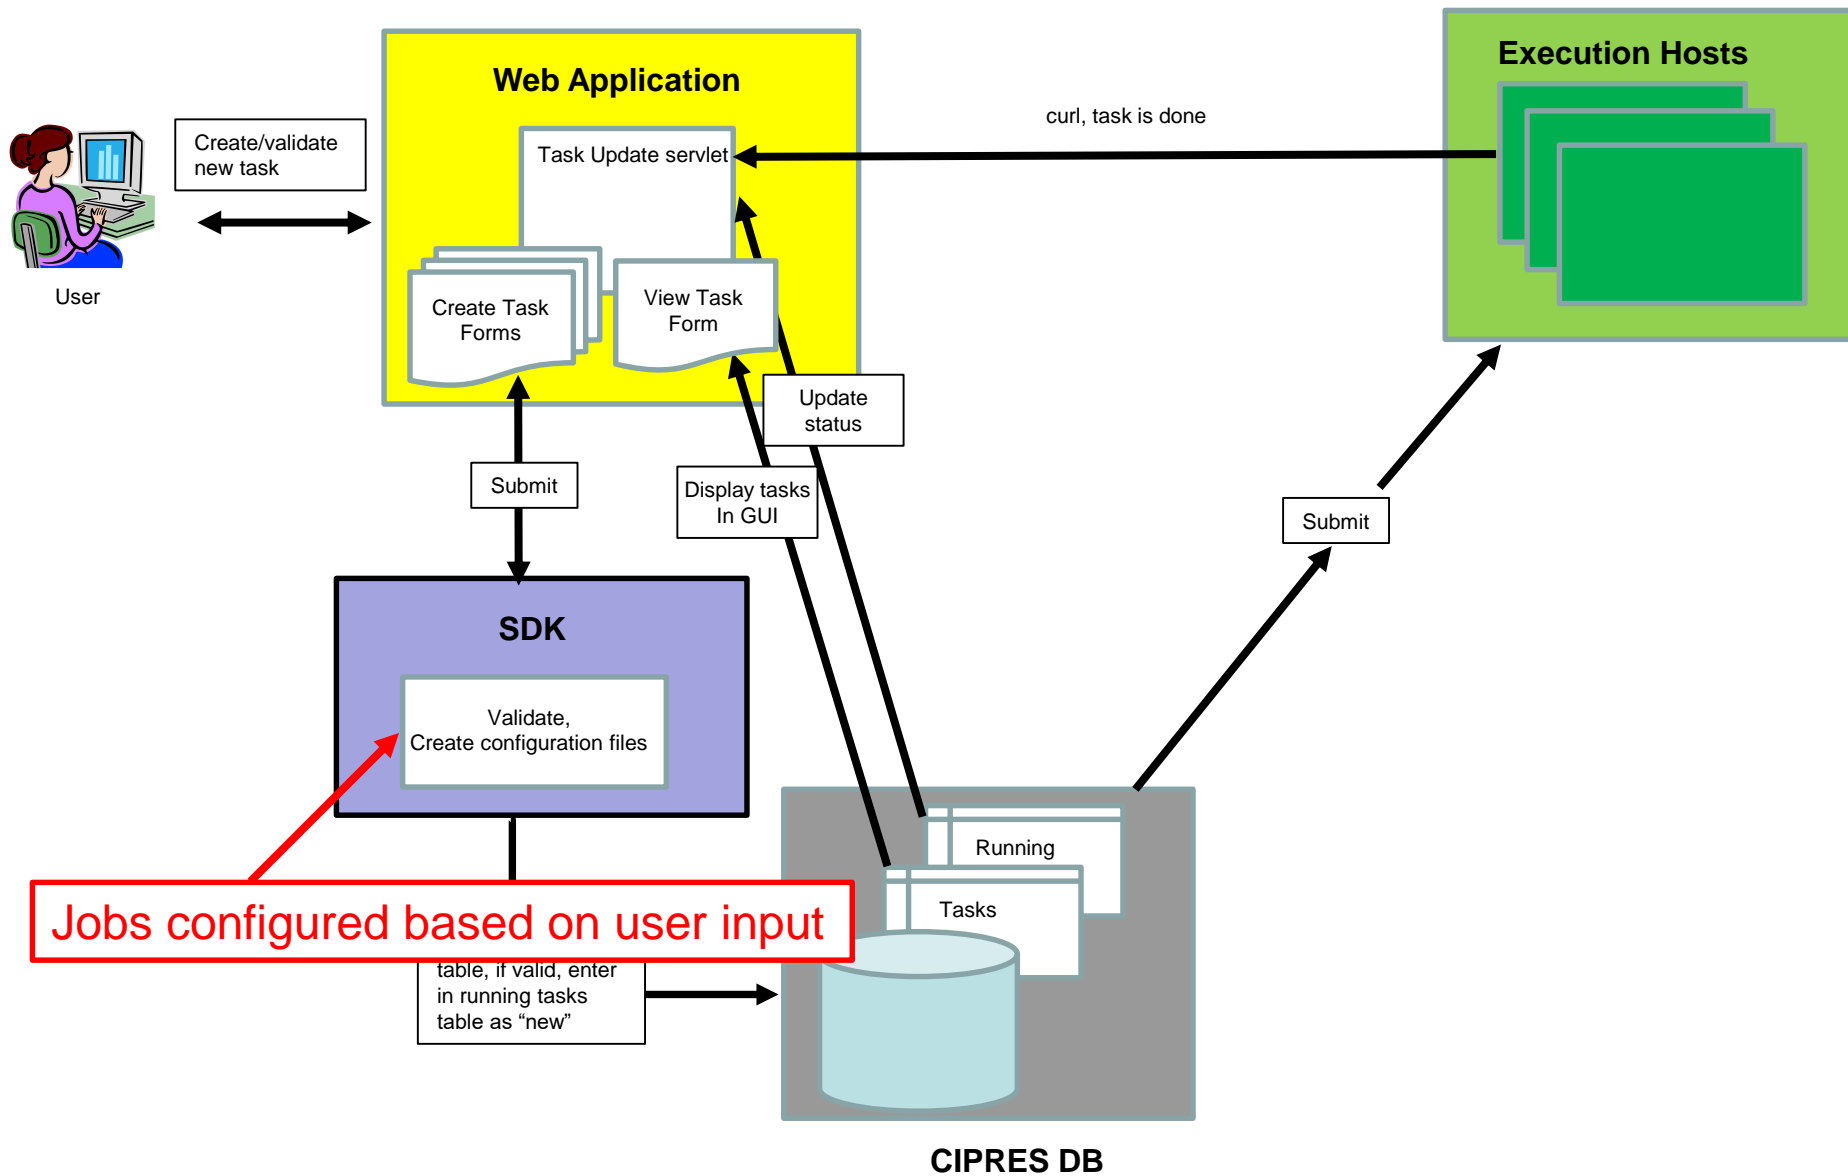
**Tasks are sent to remote execution hosts**

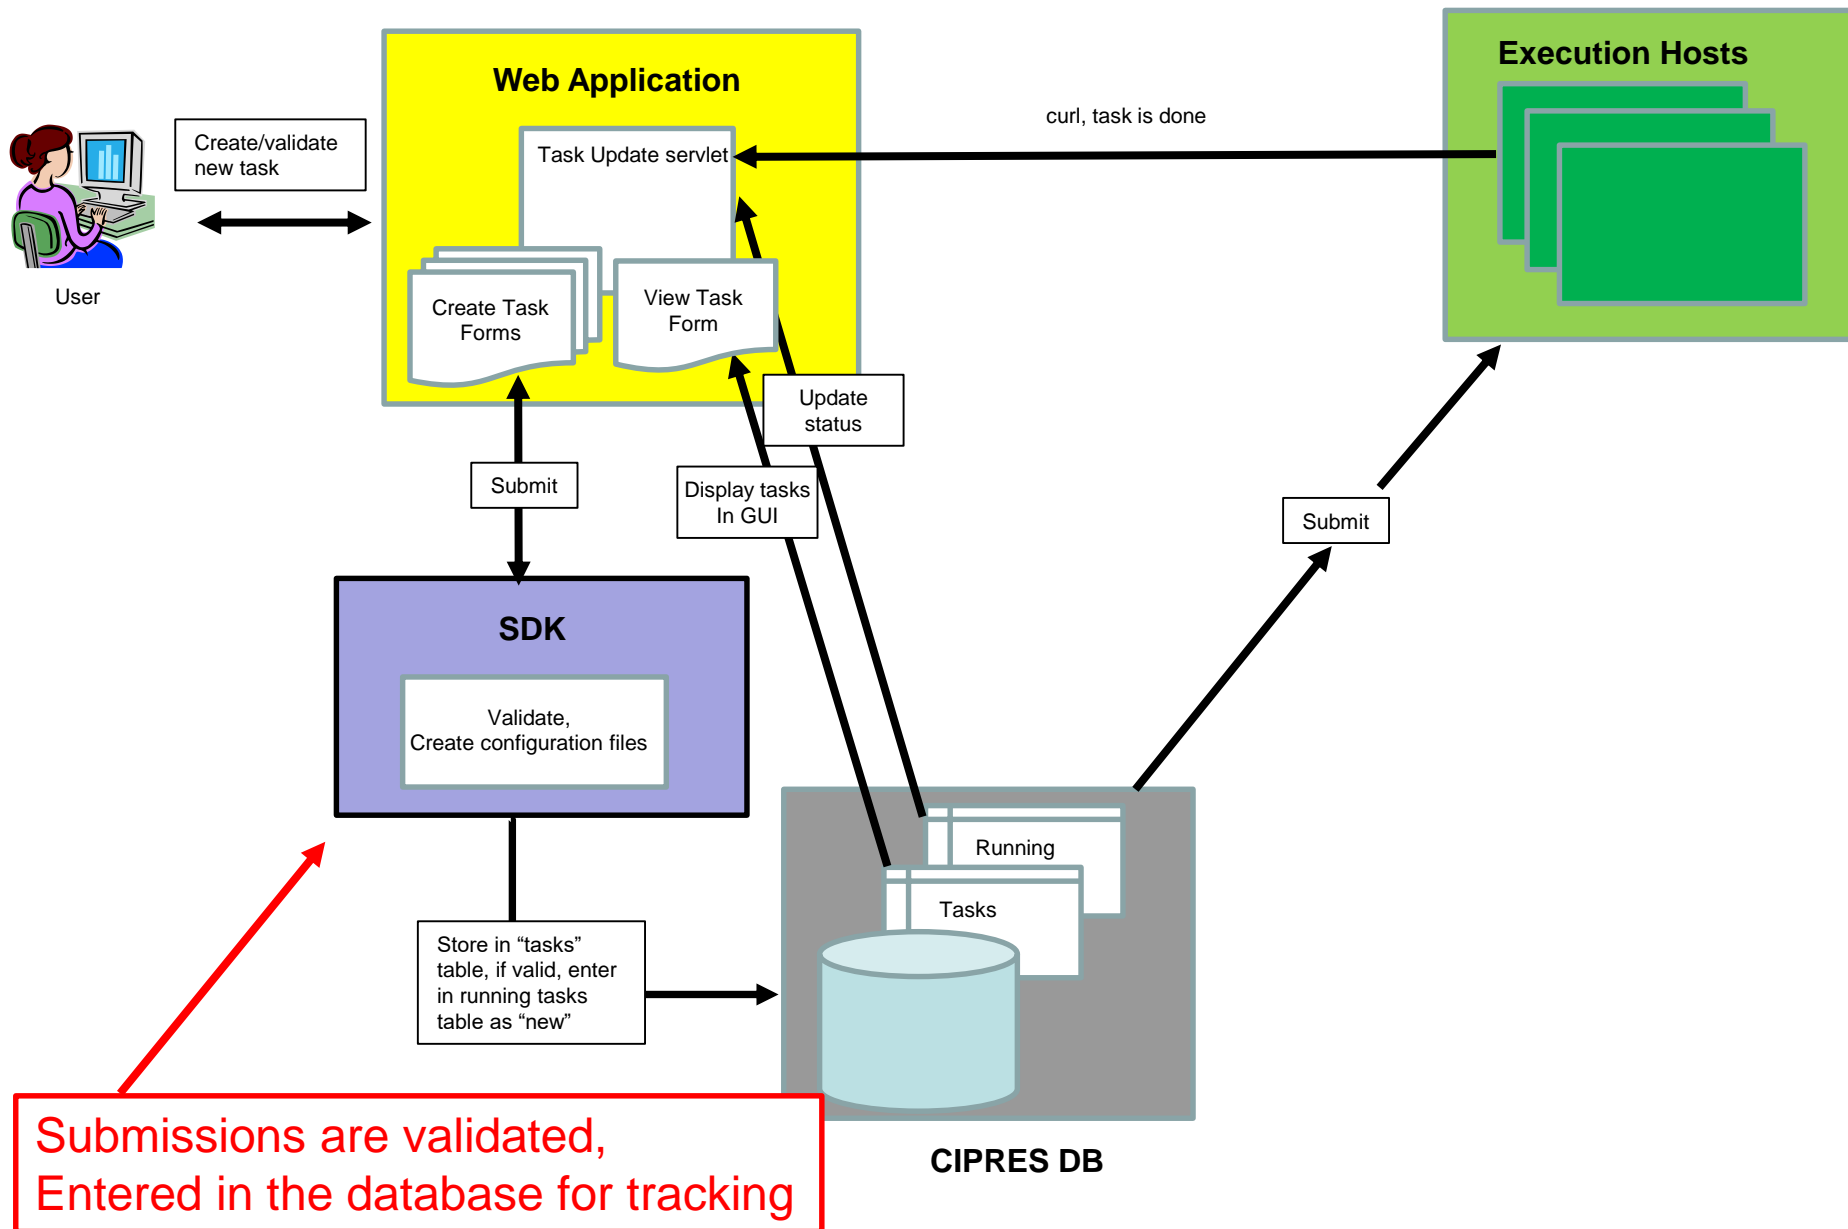


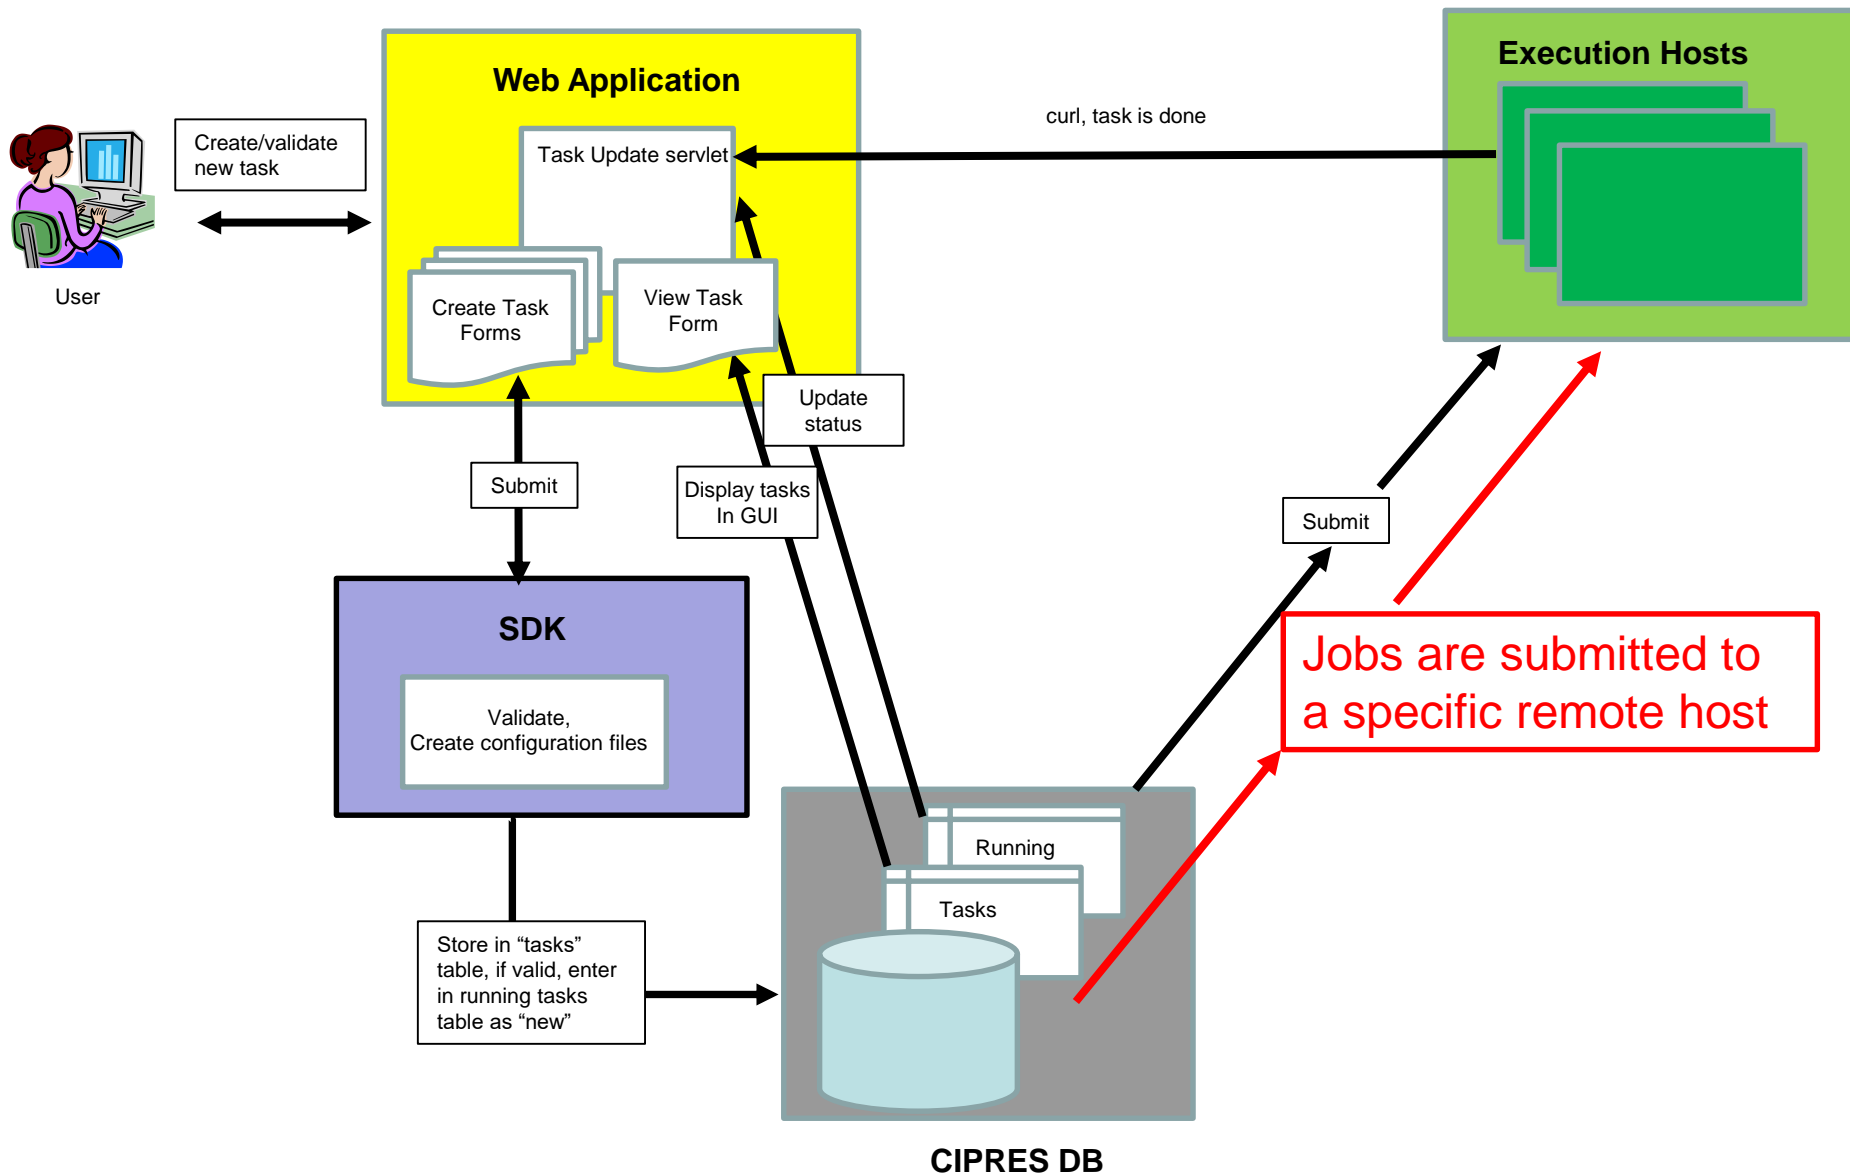


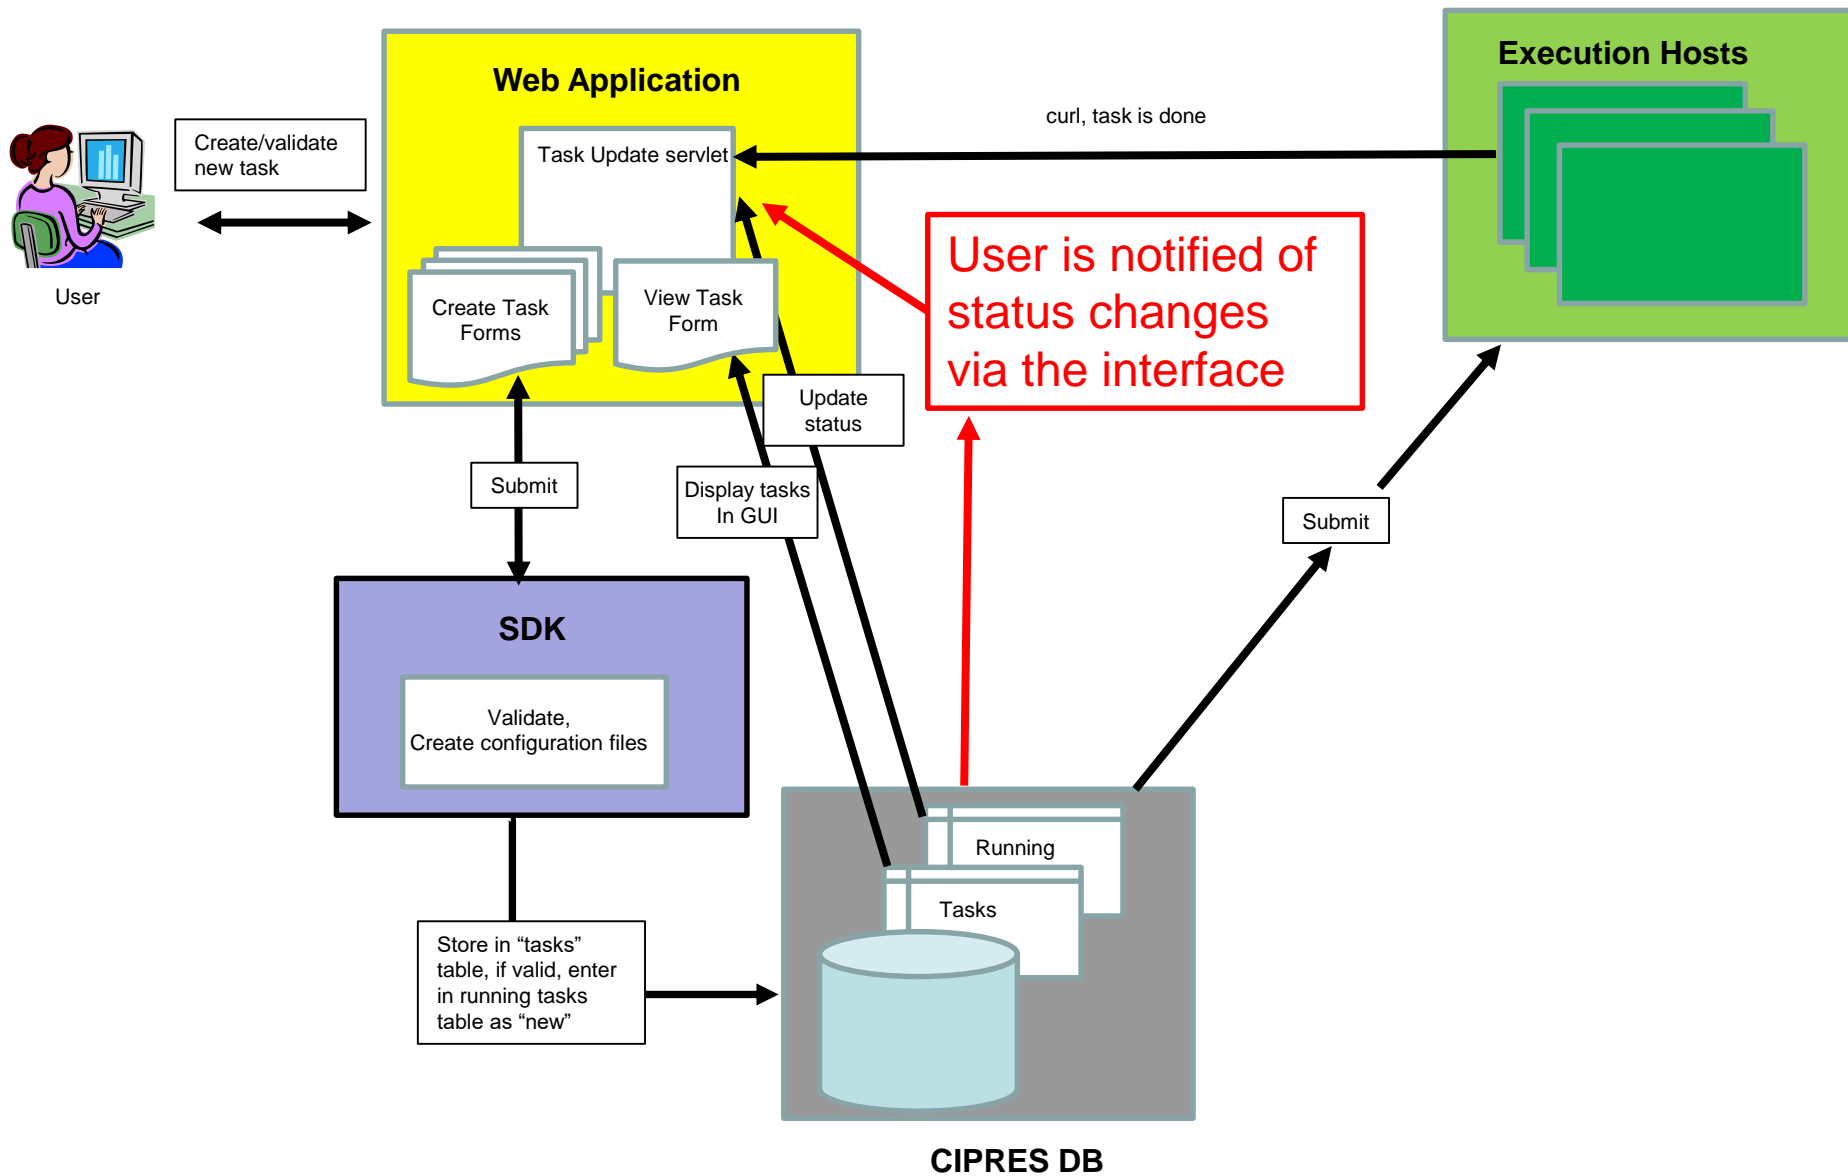














## Questions we asked early on:

- **Why do job submissions fail?**
- **When do we care a lot, a little, or not at all?**
- **What parts does CIPRES have control over**
- **What parts are not under CIPRES control?**





## Job Attrition on the CIPRES Science Gateway\*



\*March – August 2010







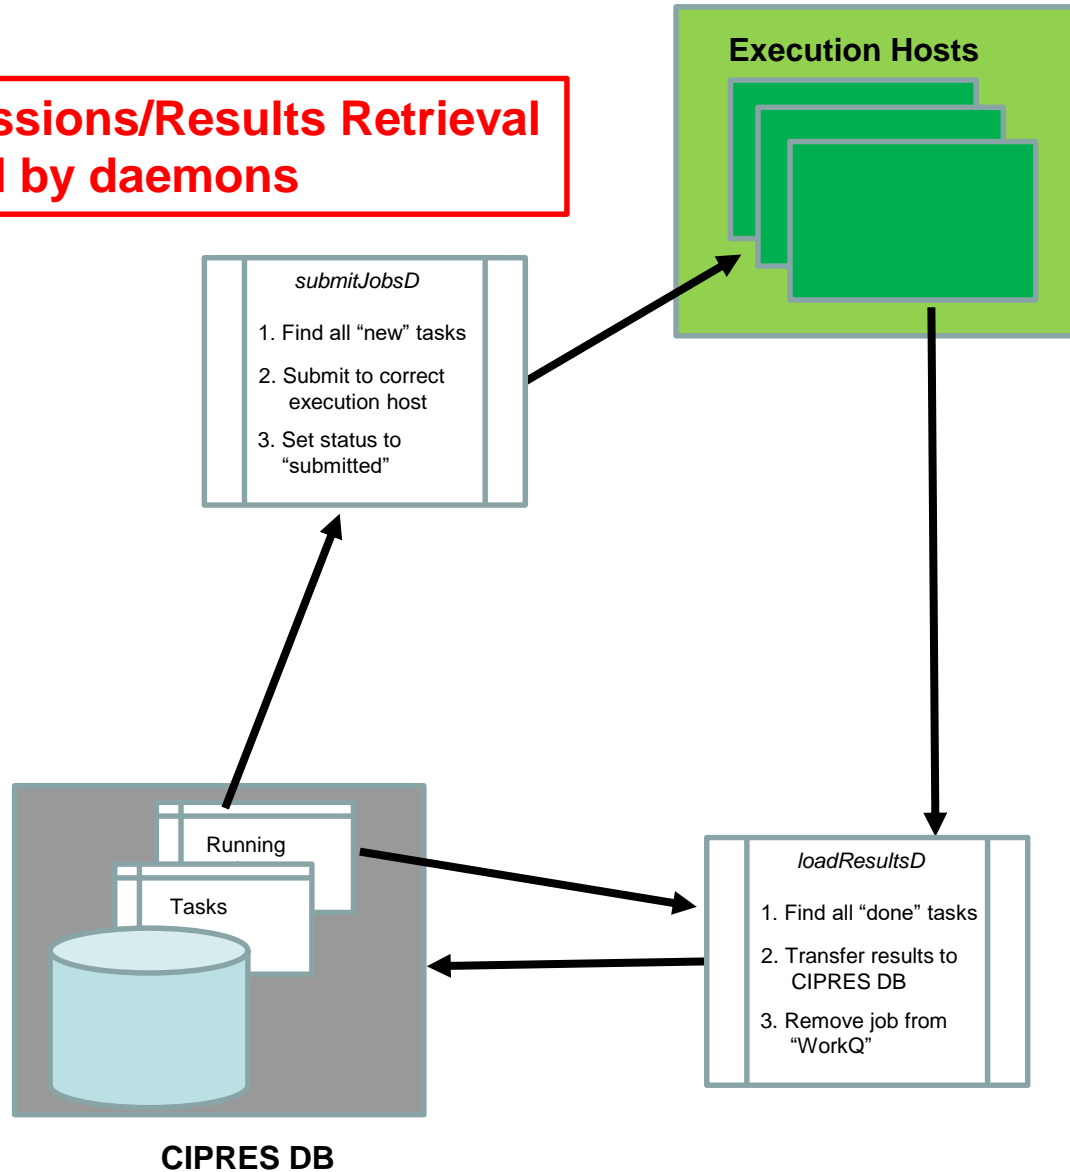
## Error Impact analysis



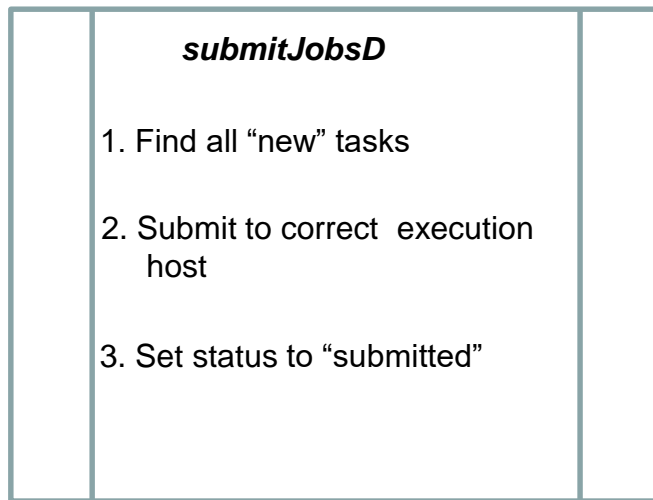
	CPU time	User	Staff
Input error (user)	0	med	med
System error	0	med	low
Communication error	high	high	high
Unknown error	?	med	low





**Job Submissions/Results Retrieval  
is managed by daemons**

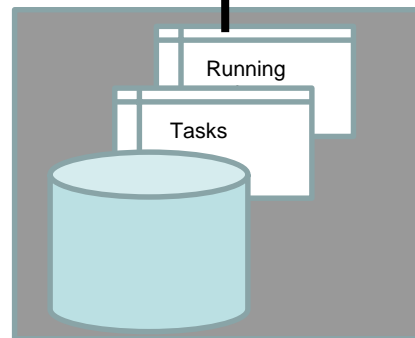
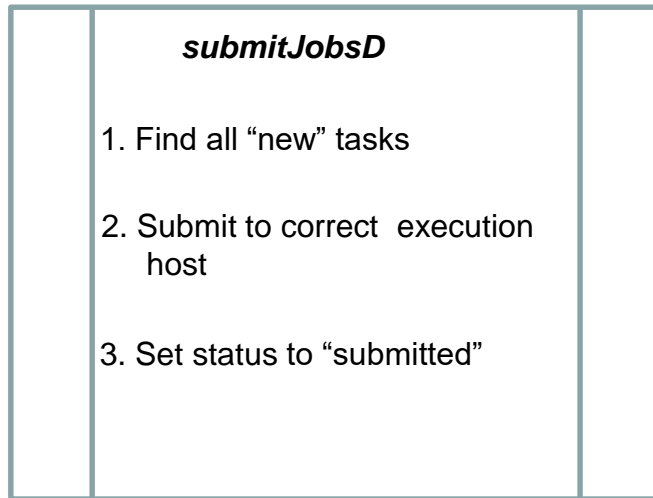




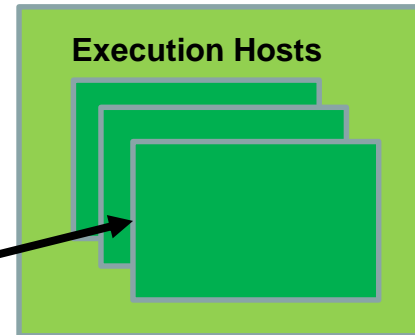
**CIPRES DB**

## **Submission daemon:**

- 1. Authenticates**
- 2. Creates working directory**
- 3. Stages input files**
- 4. Causes a script on the host to create a submit script and submit the job**



CIPRES DB



## Submission:


1. Authenticate
2. Create working directory
3. Stage input files
4. Script on submission host creates submit script

**A custom submit script is created to interact with the submit daemon for each execution host**

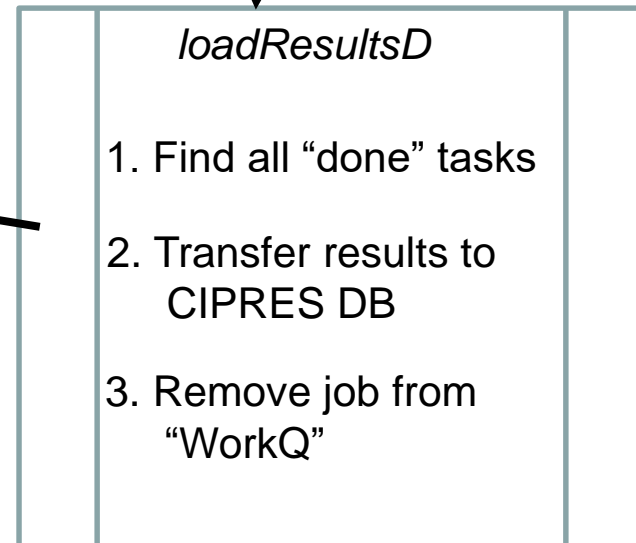



## Retrieve results daemon:

1. Communicates with a custom script on the execution host about completed tasks (e.g. `qstat -u cipres`)
2. Transfers any completed results to CIPRES DB
3. Removes the job from the WorkQ



CIPRES DB





## User errors:

**About 20% of jobs have immediate “failure” due to user error.**

**% success per code depends on whether input files are created by the user or a software package.**

**Can help users detect errors before submission, or even suggest corrections.**

**Cost of the error to CIPRES: zero, unless reported. If reported: about 1.5 fte months per year about \$15-20,000 per year.**





## System errors:

**3-4% jobs experience some form of system failure.**

***98% of system failure occurs in crisis islands;***

**Crisis island: the system is down. All submissions fail and are retried and fail again.**

***2% is sporadic scheduler transient unavailability (could be solved by re-trying). 100/250,000 jobs a year could be helped.***





## Kinds of system errors:

- **Expired or unavailable allocation**
- **Bug in the CIPRES interface causes misconfiguration (these are rare)**
- **Emergency maintenance**
- **Unavailability of the scheduler/system for unknown reasons**
- **Something unexpected in the interaction between CIPRES and the infrastructure.**







## Kinds of system errors:

- **Expired or unavailable allocation**
- **Bug in the CIPRES interface causes misconfiguration (these are rare)**
- **Emergency maintenance**
- **Unavailability of the scheduler/system for unknown reasons**
- **Something unexpected in the interaction between CIPRES and the infrastructure.**





## Features we have added:

- **ability to halt submissions from a given user account**
- **ability to monitor usage by each account automatically**
- **ability for users to track their SU consumption**
- **ability to forecast SU cost of a job for users**
- **ability for user to kill a job**
- **ability to charge to a user's personal XSEDE allocation**
- **notification of job failures mailed to staff**





## Questions we ask now:

- How can we ensure job runs are configured correctly?
- How can we detect wasteful misconfigured jobs quickly?
- What does it mean when there are no submissions?





## Features we would like to add:

- **Streamlined ability to deploy on new hosts**
- **User can determine what host a job is submitted to**
- **Route jobs to hosts based on job characteristics and host availability**
- **Automatic configuration of jobs, so users can't make mistakes**
- **Better tools for job restarting**
- **User-friendly error messaging**
- **Automated log data gathering for internal failure analysis**
- **Job submission retries on failure**

