

BV-BRC

Bacterial and Viral (BV) -
Bioinformatics Resource Center (BRC)

Monthly System Performance Metrics Report

Performance Period: June 1, 2023 – June 30, 2023

Issued to:

National Institute of Allergy and Infectious Diseases
National Institute of Health

Contract No.: 75N93019C00076

Contract Title: Bioinformatics Resource Centers for Infectious Diseases

Submission Date:

July 10, 2023

Submitted by:
Rick Stevens (contact PI)
Associate Laboratory Director
Argonne National Laboratory
Professor, Computer Science
University of Chicago
5801 South Ellis Avenue
Chicago, IL 60637-5418
630.252.3378 (phone)

BV-BRC System Performance Metrics Report

Notes: As per the recent request from NIAID, we are working with the other BRC to provide jointly agreed plots showing accumulative usage and performance data over time. We will start including them in the monthly reports, starting with the next monthly report.

This monthly system performance metrics report provides a summary of the BV-BRC system performance for the current reporting period in accordance with the Joint-BRC Common System Performance Metrics Plan developed by the BRCs and subsequently approved by NIAID.

As per the plan, each BRC will report and aggregate performance metrics for their constituent parts, *i.e.*, PATRIC and IRD/ViPR for BV-BRC. These metrics will serve as a basis for collecting quantitative measures of performance of the BRC resources to identify trends, areas that are performing well, and areas for improvement. Once the system performance plan is approved by NIAID, each BRC will submit a system performance report for their resource monthly. Annual summaries will be included in the Annual Progress Reports.

It is important to note that metrics across the two BRCs are highly dependent on the relative sizes of the respective research communities, the associated quantities and types of available data, complexity of various analysis tools, and how each of the resources delivers the data and tools to the user. Thus, cross-BRC comparisons of individual metrics are not necessarily indicative of relative usage or performance.

Common system performance metrics covering both BRCs (note that this list is subject to modification, based on feasibility of collection, changes in availability technologies, BRC website development, suggestions from NIAID program and other stakeholders, *etc.*):

Website Performance

Every month, each BRC will report the performance of the key web pages from their website, starting with the pages listed in the table below and adding new pages as they are released on the website. For each page, the average page load time will be computed based on a predefined set of pages and compared against the target page load time set as a benchmark. This will help us ensure that the performance of the individual pages and the overall website is maintained as the amount of data and usage increase with time over the life of the project. If performance of any of the pages is below the set benchmark, we will address it by performing necessary hardware or software optimizations.

- **Target page load time**

- *Definition* - Target page load time measured in seconds, set as a benchmark. The target page load times may vary for various pages depending on their complexity and amount of data they present / visualize to the user.
- *Measurement mechanism* - Manual / custom performance measurement scripts.
- *Measure* - Page load time in seconds.

- **Average page load time**

- *Definition* - Average page load time measured in seconds, after N executions. The average page load times may vary for various pages depending on their complexity and amount of data they present / visualize to the user. Hence, average load time for a web page should be compared only to the benchmark set for that page.
- *Measurement mechanism* - Manual / custom performance measurement scripts.
- *Measure* - Average page load time measured seconds, after N executions.

Table 1. BV-BRC Website Performance

Web Page	BRC Domain	Target Load Time (Seconds)	Avg. Load Time (Seconds)
Home page	BV-BRC	5	1.34
Global search – genome	BV-BRC	5	1.30
Global search – gene	BV-BRC	5	1.54
Taxonomy overview	BV-BRC	5	1.19
Genome overview	BV-BRC	5	4.37
Genome list	BV-BRC	5	2.57
Feature list	BV-BRC	5	4.46
Feature overview	BV-BRC	5	1.93
Genome browser	BV-BRC	5	2.12
Circular viewer	BV-BRC	5	2.34
SARS-CoV-2 Variants and Lineages of Concern (S-VLC) Overview	BV-BRC	5	1.21
S-VLC Lineages	BV-BRC	5	1.23
S-VLC Covariants	BV-BRC	5	1.25
S-VLC Variants	BV-BRC	5	1.39
S-VLC Genome Browser	BV-BRC	5	4.78
S-VLC Protein Structure	BV-BRC	5	1.09
S-VLC Phylogenetic Tree	BV-BRC	5	10.26
S-VLC Resources	BV-BRC	5	1.43

Notes:

1. *The IRD and ViPR resources have now been decommissioned and are no longer available. They have been removed from this summary.*

Service/Tool Performance

Both BRC analysis services and tools allow users to analyze data pulled from the respective BRC databases and their own private data, compare to other datasets, and save the results in their private workspaces. Both the BRCs will monitor and report the performance of all analysis services/tools available in their resource monthly. The performance reports will be generated based on the actual usage of these services/tools by BRC users in each month. For each analysis service, we will compute the total number of jobs submitted by users, number of jobs completed successfully, failed, average wait time for the jobs queued in the system, and average run time. Monitoring the fraction of jobs that fail and/or reported by users will allow us to identify recurring problems and address them in a timely manner to make the services more robust and reliable. The job wait time depends on the variation in the usage patterns and system load, while the run time depends heavily

on the size of the input data and the parameters selected. Monitoring these metrics will allow us to identify factors affecting the overall performance of the application services and tools and address them by performing necessary software and/or hardware scaling or optimization to meet the user expectations.

- **Analysis tasks submitted**
 - *Definition* - A breakdown of total number of analysis tasks submitted by users, summarized by service/tool, during the specified date range.
 - *Measurement mechanism* - Captured via website and server logs, which are used to tally the number.
 - *Measure* - Jobs per month, tallied by service/tool.
- **Analysis tasks completed**
 - *Definition* - A breakdown of total number of analysis tasks submitted by users and completed successfully, summarized by service/tool, during the specified date range.
 - *Measurement mechanism* - Captured via website and server logs, which are used to tally the number.
 - *Measure* - Jobs per month, tallied by service/tool.
- **Analysis tasks failed**
 - *Definition* - A breakdown of total number of analysis tasks submitted by users and failed, summarized by service/tool, during the specified date range.
 - *Measurement mechanism* - Captured via website and server logs, which are used to tally the number.
 - *Measure* - Jobs per month, tallied by service/tool.
- **Average run time by service/tool**
 - *Definition* - A breakdown of average run time for all analysis tasks submitted by users, summarized by service/tool, during the specified date range.
 - *Measurement mechanism* - Captured via website and server logs, which are used to tally the number.
 - *Measure* - Average run time measured in seconds, tallied by service/tool.
- **Input limits**
 - *Definition* - Maximum size of the input supported by a service/tool, beyond which it may degrade the performance or fail to produce results.
 - *Measurement mechanism* - Defined by requirements, design and/or testing of a service/tool.
 - *Measure* - Input size defined as number or size of the input parameters. The units can vary depending on tool/service.

Table 2. BV-BRC Tools/Services Performance Metrics

Tool/Service	BRC Domain	Jobs Submitted	Jobs Completed₁	Jobs Failed	Avg Run Time	Input Limits₂
Codon Tree	BV-BRC	731	714	17	5095	100 genomes
Comparative Systems	BV-BRC	304	285	19	518	No limit
Comprehensive Genome Analysis	BV-BRC	2707	2009	698	2920	No limit
Differential Expression	BV-BRC	7	1	6	18	No limit

FastqUtils	BV-BRC	1751	1679	71	631	No limit
Gene Tree	BV-BRC	310	208	102	274	100 long sequences, 1,000 short sequences
Genome Alignment	BV-BRC	134	125	7	1003	20 genomes
Genome Annotation	BV-BRC	4151	3934	217	257	No limit
Genome Assembly	BV-BRC	2919	2558	352	6124	No limit
Genome Comparison	BV-BRC	271	238	33	287	10 genomes
HA Subtype Numbering	BV-BRC	72	60	12	36	
Homology	BV-BRC	1914	1777	136	137	No limit
MSA	BV-BRC	511	427	84	204	No limit
MetaCATs	BV-BRC	101	101		25	10 feature groups
Metagenome Binning	BV-BRC	306	276	30	3544	No limit
Metagenomic Read Mapping	BV-BRC	146	141	3	684	No limit
Primer Design	BV-BRC	119	107	3	15	Single sequence
RNASeq Analysis	BV-BRC	1868	1257	611	2943	No limit
Sequence Submission	BV-BRC	7	5	2	743	100 sequences
Subspecies Classification	BV-BRC	320	291	29	135	No limit
Taxonomic Classification	BV-BRC	682	592	90	631	No limit
TnSeq Analysis	BV-BRC	24	12	12	1037	No limit
Variation Analysis	BV-BRC	461	416	45	1461	No limit
SARS-2 Genome Assembly and Annotation	BV-BRC	113	110	3	352	No limit

Notes:

- 1. Jobs completed will not equal submitted minus failed if the user canceled the job before completion.*
- 2. Most services do not have an imposed hard limit on the input. Job completion time is however proportional to the input size, and it may fail or hang if it is too large in some cases.*
- 3. The IRD and ViPR resources have now been decommissioned and are no longer available. They have been removed from this summary.*



Figure 1. Analysis service jobs submitted vs. completed. Months with significantly higher job submissions stemming from collaborative analyses (e.g., SARS-CoV-2 sequences) have been removed from the chart to provide more informative relative scaling from month to month.

Monitoring and resolving job failures:

The BV-BRC team monitors jobs queue daily for the number of jobs submitted, queued, and failed and addresses any issues as they occur. The primary causes of failed jobs and actions taken to address them are as follows:

1. **System errors:** Typically, if there are large number of failed jobs over a short period of time, it is due to a system-level problem, such as one or more nodes in the scheduler are temporarily unavailable, out of disk space or hung. We resolve such issues as quickly as possible, automatically resubmit failed user jobs, and notify the users who submitted the jobs.
2. **User errors:** Another common reason for failed jobs is incorrect or mal-formatted input data or parameters provided by the user. We address such errors by informing users about the correct input formats by pointing them to the relevant documentation and/or examples or working with them to resolve the issue. We identify any recurring patterns and update user documentation and/or job submission forms as needed to avoid such errors in future.
3. **Bug in the application:** Sometimes job failure can result from an unknown bug in a specific application, triggered by the combination of the format or size of the user input data and/or parameters provided by the user. We address this by identifying the cause of failure and then fixing it as soon as possible.

In future, we will establish limits on failure rate that will trigger further investigation of the root causes of the problem and potential solutions to address the and include them in the monthly report.

Database / Data API Performance

Both the BRCs will monitor database performance using predefined search and retrieval queries for various data types, measure average response time in seconds, and report it on a monthly basis. These database queries will capture the most common data queries used by various web pages and tools on the BRC websites as well as user queries used to download large amounts of data in batch mode using the data API, web services, or Command Line Interface (CLI). For each query, the average response time will be compared to the set benchmark. This will help us ensure that the performance of individual data queries as well as the overall database meets the performance benchmarks as well as user expectations. If the performance of any query does not meet the benchmark, we will address it by performing necessary database, query, or hardware optimizations.

- **Target response time**

- *Definition* - Target response time measured in seconds, set as a benchmark. The target response times may vary for various queries depending on the complexity of the query and amount of data retrieved.
- *Measurement mechanism* - Manual / custom performance measurement scripts.
- *Measure* - Page load time in seconds.

- **Average response time**

- *Definition* - Average response time measured in seconds, after N executions. The average response times may vary for various pages depending on the complexity of the query and amount of data retrieved. Hence, average load time for a web page should be compared only to the benchmark set for that page.
- *Measurement mechanism* - Manual / custom performance measurement scripts.
- *Measure* - Average response time measured seconds, after N executions.
- *Note: For IRD/ViPR, the average response times reported are representative of both resources, as they both are part of the same single application and database.*

Table 3. BV-BRC Database / Data API Performance

Database Query	BRC Domain	Target Response Time (Seconds)	Avg Response Time (Seconds)
Genome	BV-BRC	5	0.09
Genome sequence	BV-BRC	5	2.10
Genome feature	BV-BRC	5	0.85
Genome AMR	BV-BRC	5	0.06
Feature sequence	BV-BRC	5	0.33
Specialty gene	BV-BRC	5	0.05
Pathway	BV-BRC	5	0.84
Subsystem	BV-BRC	5	0.75
Taxonomy	BV-BRC	5	0.09
Transcriptomics Experiment	BV-BRC	5	0.04
Transcriptomics Sample	BV-BRC	5	0.09
Transcriptomics Genes	BV-BRC	5	0.10

PPI	BV-BRC	5	0.66
SARS-CoV-2 Lineages	BV-BRC	5	0.37
SARS-CoV-2 Variants	BV-BRC	5	0.42