

BV-BRC

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

Monthly System Performance Metrics Report

Performance Period: September 1, 2022 – September 30, 2022

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:

October 10, 2022

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)
630.252.6333 (fax)

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.

- *Note: For IRD/ViPR, the average page load times reported are representative of both resources, as they both are part of the same single application and database.*

Table 1. BV-BRC Website Performance

Web Page	BRC Domain	Target Load Time (Seconds)	Avg. Load Time (Seconds)
PATRIC home page	PATRIC	5	0.94
Global search – genome	PATRIC	5	2.35
Global search – gene	PATRIC	5	2.12
Taxonomy overview	PATRIC	5	1.81
Genome overview	PATRIC	5	5.95
Genome list	PATRIC	5	4.23
Feature list	PATRIC	5	1.21
Feature overview	PATRIC	5	3.61
Genome browser	PATRIC	5	2.54
Circular viewer	PATRIC	5	3.58
Home Page	IRD/ViPR	5	1.22
Genome Search	IRD/ViPR	10	3.47
Protein Search	IRD/ViPR	10	2.84
Strain Search	IRD/ViPR	10	3.21
Epitope Search	IRD/ViPR	10	4.79
Surveillance Search	IRD/ViPR	10	3.8
SFVT Search	IRD/ViPR	10	0.915
Strain Details Page	IRD/ViPR	5	1.42
Protein Details Page	IRD/ViPR	5	1.12
Protein Structure Search	IRD/ViPR	5	3.5
SARS-CoV-2 Variants and Lineages of Concern (S-VLC) Overview	BV-BRC	5	0.95
S-VLC Lineages	BV-BRC	5	0.87
S-VLC Covariants	BV-BRC	5	1.08
S-VLC Variants	BV-BRC	5	1.30
S-VLC Genome Browser	BV-BRC	5	4.47

S-VLC Protein Structure	BV-BRC	5	0.90
S-VLC Phylogenetic Tree	BV-BRC	5	6.79
S-VLC Resources	BV-BRC	5	0.91

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/PATRIC	553	537	16	6,284	100 genomes
Comparative Systems	BV-BRC	157	130	27	666	No limit
Comprehensive Genome Analysis	BV-BRC/PATRIC	1,606	1,399	207	4,354	No limit
Differential Expression	BV-BRC/PATRIC	8	6	2	23	No limit
FastqUtils	BV-BRC/PATRIC	333	323	10	1,091	No limit
Gene Tree	BV-BRC	35	28	7	405	100 long sequences, 1,000 short sequences
Genome Alignment	BV-BRC/PATRIC	178	163	15	644	20 genomes
Genome Annotation	BV-BRC/PATRIC	4,340	4,116	224	523	No limit
Genome Assembly	BV-BRC/PATRIC	6,791	4,971	1,820	3,555	No limit
Genome Comparison	BV-BRC/PATRIC	407	360	47	430	10 genomes
Homology	BV-BRC	1,038	972	66	659	No limit
MSA	BV-BRC	204	168	36	67	No limit
MetaCATs	BV-BRC	57	50	7	21	
Metagenome Binning	BV-BRC/PATRIC	222	174	48	7,231	No limit
Metagenomic Read Mapping	BV-BRC/PATRIC	134	131	3	1,998	No limit
Primer Design	BV-BRC	42	42	0	12	Single sequence
RNASeq Analysis	BV-BRC/PATRIC	88	66	22	5,544	No limit
Taxonomic Classification	BV-BRC/PATRIC	399	377	22	1,756	No limit
TnSeq Analysis	BV-BRC/PATRIC	70	43	27	2,393	No limit
Variation Analysis	BV-BRC/PATRIC	469	377	92	3,767	No limit
Alignment Viewer	IRD/ViPR	68	64	4	1,444	No limit

Antiviral-Resistance-Risk	IRD/ViPR	15	15	0	4	No limit
BLAST	IRD/ViPR	259	243	16	2,301	No limit
Enrichment	IRD/ViPR	0	0	0	0	No limit
Genotype-Recombination	IRD/ViPR	9	6	3	7,692	No limit
H1-Clade Classifier	IRD/ViPR	113	108	5	1,329	No limit
H1N1-classifier	IRD/ViPR	12	8	4	8,118	No limit
H5N1-classifier	IRD/ViPR	93	83	10	2,627	No limit
Ha Numbering	IRD/ViPR	92	87	5	1,729	No limit
MGC	IRD/ViPR	12	11	1	2,515	No limit
MSA	IRD/ViPR	605	513	92	8,199	No limit
Mutation-analysis	IRD/ViPR	15	13	2	3,326	No limit
Primer3	IRD/ViPR	53	47	6	2,971	No limit
Read-seq	IRD/ViPR	29	29	0	7	No limit
Rva Genotyper	IRD/ViPR	904	857	47	1,823	No limit
Short-seqsearch	IRD/ViPR	16	5	11	16,817	No limit
SNP-analysis	IRD/ViPR	450	387	63	13,497	No limit
Surveillance-data-mapping	IRD/ViPR	6	6	0	2	No limit
Tbl-formatter	IRD/ViPR	3	3	0	1	No limit
Tree	IRD/ViPR	334	300	34	7,704	No limit
VIGOR Annotator	IRD/ViPR	42	40	2	1,450	No limit
SARS-2 Genome Assembly and Annotation	BV-BRC	77	69	8	690	No limit

1. Jobs completed will not equal submitted minus failed if the user canceled the job before completion.
2. These 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.




Figure 1. Analysis service jobs submitted vs. completed. February, March, and May 2021 had significantly higher job submissions stemming from collaborative SARS-CoV-2 sequence analysis.

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 ₁	Target Response Time (Seconds)	Avg Response Time (Seconds)
Genome	BV-BRC/PATRIC	5	0.06
Genome sequence	BV-BRC/PATRIC	5	1.21
Genome feature	BV-BRC/PATRIC	5	0.58
Genome AMR	BV-BRC/PATRIC	5	0.04
Feature sequence	BV-BRC/PATRIC	5	0.18
Specialty gene	BV-BRC/PATRIC	5	0.04
Pathway	BV-BRC/PATRIC	5	1.09
Subsystem	BV-BRC/PATRIC	5	0.95
Taxonomy	BV-BRC/PATRIC	5	0.03
Transcriptomics Experiment	BV-BRC/PATRIC	5	0.03
Transcriptomics Sample	BV-BRC/PATRIC	5	0.09
Transcriptomics Genes	BV-BRC/PATRIC	5	0.88
PPI	BV-BRC/PATRIC	5	0.59
Sequence & Strain	IRD/ViPR	5	0.05

Animal Surveillance	IRD/ViPR	5	0.01
3D Protein Structures	IRD/ViPR	5	0.001
Human Surveillance	IRD/ViPR	5	0.01
Serology Experiments	IRD/ViPR	5	0.02
Antiviral Drugs	IRD/ViPR	5	0.01
Host Factor Data	IRD/ViPR	5	0.01
Immune Epitopes	IRD/ViPR	5	0.02
Phenotypes	IRD/ViPR	5	0.02
PCR Primer Probe Data	IRD/ViPR	5	0.001
Sequence Feature Variant Types	IRD/ViPR	5	0.01
Human Clinical Studies and Lab Experiments	IRD/ViPR	5	0.001
ORFeome Plasmid Data	IRD/ViPR	5	0.01
SARS-CoV-2 Lineages	BV-BRC	5	0.27
SARS-CoV-2 Variants	BV-BRC	5	0.31

1. BV-BRC and PATRIC use the same backend database and API.