



Bionano Compute On Demand v1.2 Release Notes

Document Number: 30293

Document Revision: N

Table of Contents

Legal Notice	3
Revision History	4
Bionano Compute On Demand	4
Introduction	4
Audience	4
Compatibility	5
Internet Access	5
Capacity and Performance	5
How to Enable the Service	6
Compute On Demand Operations	9
Tokens	9
Features / Benefits	10
Terms and Conditions	13
High Performance Storage Architecture	14
Tickets	14
Token Ranges	15
Technical Assistance	16

Legal Notice

For Research Use Only. Not for use in diagnostic procedures.

This material is protected by United States Copyright Law and International Treaties. Unauthorized use of this material is prohibited. No part of the publication may be copied, reproduced, distributed, translated, reverse-engineered or transmitted in any form or by any media, or by any means, whether now known or unknown, without the express prior permission in writing from Bionano Genomics. Copying, under the law, includes translating into another language or format. The technical data contained herein is intended for ultimate destinations permitted by U.S. law. Diversion contrary to U. S. law prohibited. This publication represents the latest information available at the time of release. Due to continuous efforts to improve the product, technical changes may occur that are not reflected in this document. Bionano Genomics reserves the right to make changes in specifications and other information contained in this publication at any time and without prior notice. Please contact Bionano Genomics Customer Support for the latest information.

BIONANO GENOMICS DISCLAIMS ALL WARRANTIES WITH RESPECT TO THIS DOCUMENT, EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO THOSE OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. TO THE FULLEST EXTENT ALLOWED BY LAW, IN NO EVENT SHALL BIONANO GENOMICS BE LIABLE, WHETHER IN CONTRACT, TORT, WARRANTY, OR UNDER ANY STATUTE OR ON ANY OTHER BASIS FOR SPECIAL, INCIDENTAL, INDIRECT, PUNITIVE, MULTIPLE OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH OR ARISING FROM THIS DOCUMENT, INCLUDING BUT NOT LIMITED TO THE USE THEREOF, WHETHER OR NOT FORESEEABLE AND WHETHER OR NOT BIONANO GENOMICS IS ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Patents

Products of Bionano Genomics® may be covered by one or more U.S. or foreign patents.

Trademarks

The Bionano Genomics logo and names of Bionano Genomics products or services are registered trademarks or trademarks owned by Bionano Genomics in the United States and certain other countries.

Bionano Genomics®, Irys®, IrysView®, IrysChip®, IrysPrep®, IrysSolve®, Saphyr®, Saphyr Chip®, Bionano Access®, and Bionano EnFocus™ trademarks of Bionano Genomics, Inc. All other trademarks are the sole property of their respective owners.

No license to use any trademarks of Bionano Genomics is given or implied. Users are not permitted to use these trademarks without the prior written consent of Bionano Genomics. The use of these trademarks or any other materials, except as permitted herein, is expressly prohibited and may be in violation of federal or other applicable laws.

© Copyright 2021 Bionano Genomics, Inc. All rights reserved.

Revision History

Revision	Notes
A	Initial release of document.
K	Updates for Compute On Demand version 1.1
L	Updates for Compute On Demand version 1.2
M	Non-human token ranges updated.
N	Updates for Compute On Demand version 1.2.1

Bionano Compute On Demand

This document describes the v1.2.1 release of Bionano Compute On Demand. In this document we will provide an overview of the service and how to access it. Should you have any questions please contact support@bionanogenomics.com.

Introduction

Bionano Compute On Demand is a hosted service for performing Bionano Solve operations. It is accessible through your Bionano Access web server. Once enabled (the process is described below), users can employ cloud resources to execute various analyses of Bionano data using the Bionano Solve pipelines. Bionano Compute On Demand is a pay-per-use service. Customers use tokens to pay for their jobs. Each job is executed on its own independent compute environment, and the results are returned to your Bionano Access system. Once the results have been successfully downloaded to your Bionano Access, they are deleted from the Bionano Compute On Demand system.

Audience

Bionano Compute On Demand is a useful product for all users of Bionano Access. Here are some scenarios that would benefit from this offering:

- You have Bionano data but no compute resources of your own.
- It is not possible for your organization to host compute servers.
- You have an operation you need to expedite, and your compute cluster is busy.
- You require capacity beyond your existing compute servers for a short period of time.
- You don't need compute resources all the time.
- You process plant or animal genomes larger than 5 Gbp.

Compatibility

Bionano Compute On Demand v1.2 is compatible with Bionano Access 1.5 or greater. The version of Bionano Solve used for compute on Bionano's Compute On Demand service will depend on which version of Bionano Access is used. Bionano Access 1.6 is designed to work with Bionano Solve 3.6. This way, the Bionano Compute On Demand results are consistent with any compute jobs you might perform locally. It also allows users to control when they upgrade to the latest pipeline.

Internet Access

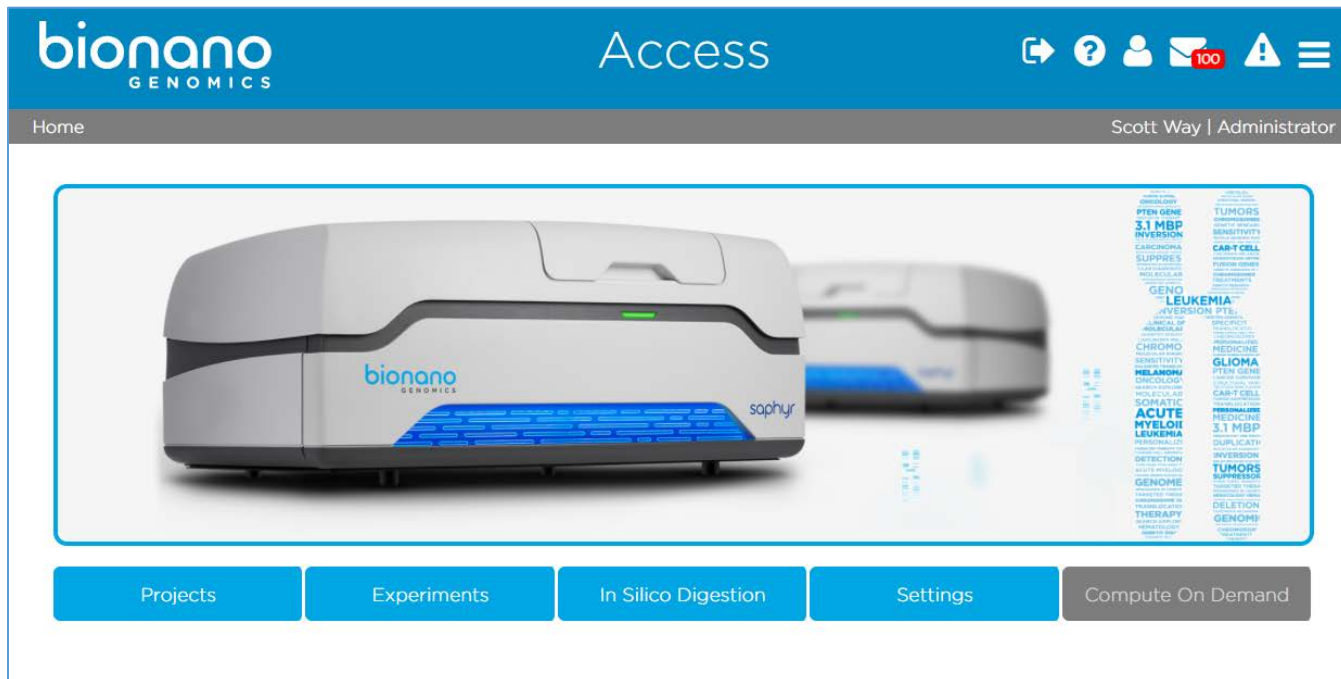
To interface with the Bionano Compute On Demand service, your Bionano Access server must have internet access. We recommend that you have 20 Mbps upload and download speed to accommodate large file uploads and downloads. Slower speeds will work as long as your connection is stable, but file transfers could take considerable time depending on the speed of your connection. A typical human *de novo* assembly result file could be 5-10 GB.

Capacity and Performance

There are limitations on the size of input data for operations the Bionano Compute On Demand service will allow. For this release, *de novo* assembly operations accept BNX files with less than 4 Tbp of DNA longer than 150 kbp. Rare variant analysis operations accept BNX files with less than 6 Tbp of DNA longer than 150 kbp. When submitting a job with more data than these limits, Bionano Access will show a warning and not allow that data to be processed. The data can be down-sampled using Bionano Access. Please review the various Bionano Theory of Operations for the coverage recommended for the different operations.

How to Enable the Service

To enable Compute on Demand you must be an Administrator. Your role should appear in the breadcrumbs bar in the header. From the home page, click the 'Settings' button.



From the settings screen, click the 'System Settings' button.



Then on the system settings page, click the slider in the 'Compute On Demand and Saphyr Assure' panel.

The screenshot shows a configuration page with two main sections: 'System Level Services' and 'User Account Settings'.

System Level Services:


- Compute On Demand and Saphyr Assure:** A yellow header. Description: "The Compute On Demand service provides pay per use elastic computing resources to supplement and accelerate your analysis needs. Together with the Saphyr Assure service you can monitor the health of your entire workflow. This is the recommended option." The toggle is currently turned **On**. A blue 'Test Connection' button is present. A link for 'Terms and Conditions' is at the bottom.
- Saphyr Assure Only:** A blue header. Description: "Use this option if you want to ensure optimal performance of your instrument, but are not interested in the Compute On Demand service." The toggle is currently turned **Off**. A link for 'Terms and Conditions' is at the bottom.
- Maintenance Mode:** A grey header. Description: "In Maintenance Mode the system can only be accessed by users with the Administrator Role. This is useful during and after updates to ensure proper operation of the system before releasing it back to users." The toggle is currently turned **Off**.

User Account Settings:

- Idle User Session Expiration [minutes]:** 999
- Maximum User Concurrent Sessions:** 1
- Maximum Login Attempts:** 3
- Password Retention:**
 - Password Retention Checks:** Off On
 - Password Expiration [days]:** 90
 - Previous Password Retention [months]:** 3
- Password Settings:**
 - Minimum Characters:** 0
 - Minimum Numerals:** 0
 - Minimum Special Characters:** 0
 - Minimum Case Changes:** 0

At the bottom right of the User Account Settings section are 'Submit' and 'Cancel' buttons.

If this is the first time you are enabling the service, you will be prompted for some information and asked to agree to the Compute on Demand Terms and Conditions. Provide the information requested then click 'I Accept'. You can view the Terms and Conditions anytime by clicking the Terms and Conditions link.



Bionano Compute On Demand Activation

Bionano Compute On Demand Registration Information

Organization / Company Name

Region
North America

Terms and Conditions

The following spell-out the terms of use of Bionano Compute On Demand. By selecting 'I Accept' you agree to the following:

- You may only use Bionano Compute On Demand to process data that is owned, licensed or lawfully obtained by you.
- You must comply with the current Bionano technical documentation applicable to the data you send to the Bionano Compute On Demand. Noncompliance to this technical documentation may result in failure of operations that would not be reimbursable by Bionano.
- Data will be transferred from your local server to a cloud resource for the purpose of computing in the cloud.
- Upon successful completion, data are returned to this Access Server and are deleted from the cloud resource.
- When operations fail in the cloud, data are retained temporarily within the Bionano Compute On Demand servers. You may authorize Bionano Support to view this data for the sole purpose of troubleshooting the failure. Whether or not you contact Bionano, this temporary data will be deleted pursuant to Bionano's procedures in effect at that time.
- Tokens provided from Bionano are required to run operations in the Bionano Compute On Demand. You will act ethically in procuring and using tokens. You will not attempt to misuse or otherwise attempt to subvert the use of tokens.
- Metrics regarding Bionano Compute On Demand operations will be collected to support billing and product improvement. These metrics include information on the status of your system and are not shared with third parties. Bionano does not collect any protected health information nor your actual run data or pipeline results.
- You agree to receive notification email before any Compute On Demand system maintenance and upgrade.
- Bionano, at its sole discretion and without prior notice, may terminate your access to the Bionano Compute On Demand if you are not in compliance with the terms of use or otherwise behave in an illegal or unethical way.

If you are not ready to complete the opt-in dialog, you can click the 'Cancel' button, but Administrators will continue to be prompted each time they login until they complete the opt-in dialog or until the 'BCOD' setting in your configuration file is removed. Once the opt-in has been completed, the dialog will no longer appear. You will not be able to change your selections later, so complete the form carefully.

The first prompt is for your Organization / Company Name. We recommend setting this value to the domain in your email. For example if your email was jsmith@bionanogenomics.com, we would recommend setting your organization to 'bionanogenomics'. Do not use a generic name like 'Instrument' or 'Saphyr' for your organization. The organization name is used to identify jobs launched from your instance of Bionano Access. Each Bionano Access server must have a unique organization name. Your organization name will not be accepted if it is not unique.

The second prompt is for your computing region. We currently offer three computing end points. There is one for North America, one for Europe, and one for Germany. Additional regions will be added in the future based on demand. The region is where the servers that will execute your compute jobs will reside. You are not restricted to which region you select, but selecting a region closest where your system resides is best for security and performance reasons. We do not recommend selecting Germany unless you are under restriction to operate only in that region.

Compute On Demand Operations

Once the opt-in process described above has been completed, Bionano Compute On Demand options will appear in the Project Browser. In the Project Browser, the options panel will display different operations you can perform based on the object type selected. When Bionano Compute On Demand is enabled, there will be an additional panel visible below the Options panel (image below). The Bionano Compute On Demand options are listed separately since they require tokens. Like the existing Options panel, the hyperlinks on the Compute on Demand panel will change depending on the type of object selected.



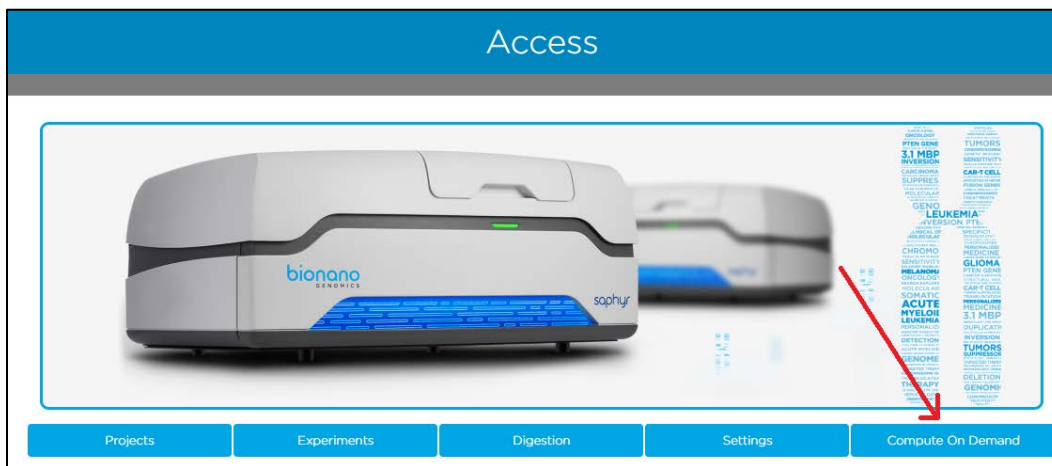
Tokens

When performing a Bionano Solve operation locally, you first will be prompted for the required inputs. When performing a Bionano Compute On Demand operation, the inputs are the same as the inputs with one addition, the token cost (image below). To proceed, you must have sufficient tokens and you must approve the token cost. The token cost estimates will be a range. In the example below, the cost estimate is between 4 and 6 tokens. When you approve the operation the maximum token cost (6 in this example) will be deducted from your token balance. When the job completes, if the cost was less than the maximum token estimate, you would be refunded the difference. For example, if the job below were to cost 4 tokens upon completion, you would receive a message informing you that 2 tokens have been returned to your token balance. You are guaranteed that the cost will not exceed the maximum token cost quoted.

Token Cost

Thank you for requesting a de novo Assembly. We performed a molecule quality report (MQR) using your data to gauge the volume and quality of the data that will be used. Based on this information it will cost between **9 and 14 tokens** to produce your assembly. By clicking the **Submit** button you acknowledge this cost and authorize us to reserve the maximum tokens. Once your operation has completed Bionano Compute On Demand will reconcile the actual cost against the tokens reserved. Your current token balance is **37**. Click **Approve and Submit** to continue.

Tokens to execute jobs on our Bionano Compute On Demand service can be acquired from your designated sales person or by contacting orders@bionanogenomics.com. When you purchase tokens, you will receive one or more voucher codes. You can redeem a voucher by clicking on the 'Compute On Demand' button on the main screen in Access (image below).



When a voucher is redeemed, all the tokens for that voucher are granted to the account for the user who is logged in. Any user who owns tokens can transfer them to other user accounts (except those with Read-Only access) on the same server. Tokens currently cannot be transferred between Bionano Access servers. Tokens cannot be taken they can only be given. Read-Only accounts cannot have tokens, because they cannot perform operations.

Base Features / Benefits

Below is a brief description of new features and benefits of this release. For more detailed information on these new features, please refer to our user manuals and video tutorials.

Feature	Description
Tokens	Token vouchers are used instead of direct payment methods for security reasons. By using tokens, no credit card information or payment methods are accessible through your Bionano Access server or Bionano Compute On Demand. Tokens are only available through the Bionano sales team or at orders@bionanogenomics.com . Depending on demand we may offer a self-serve voucher site in the future.
Estimate Ranges	Job token costs are presented to you in the form of a range. We believe you deserve to know the potential cost up front, but we also want to only charge you based on the actual cost. This way you get the best of both worlds. You are not charged more than the maximum stated, but if the job costs less, we will refund you the difference. Due to the complexity of the data we process and the quality of your data, it is not always possible to predict the exact cost up front, hence we provide a range. We are continuously working on narrowing the range we provide.
Data Security	Data security is always a primary concern when you consider a hosted solution. A lot of care was given to our system design in order to provide the best assurances your data is protected. For more detailed information on data security please refer to our separate Data Security Guideline document (30292). It is important to note that data is returned to your local Bionano Access server upon successful completion, and Bionano Compute On Demand does not retain protected data. When operations fail in Bionano Compute On Demand, data are retained temporarily within the Bionano Compute On Demand servers. You may authorize Bionano Support to view this data for the sole purpose of troubleshooting the failure. Whether or not you contact Bionano, this temporary data will be deleted pursuant to Bionano's procedures in effect at that time.
Bionano Access Integration	The Bionano Compute On Demand solution is integrated with Bionano Access. Results from Bionano Compute On Demand will function in Bionano Access the same as those from local compute resources. Users already familiar with how to use Bionano Access will understand how to use Bionano Compute On Demand operations quickly, because they are nearly identical.
Encryption	All data transfers and communication to and from Bionano Compute On Demand are encrypted.
Agnostic Solution	Our cloud solution is provider agnostic. Our solution has been designed to leverage whichever cloud provider has the best value whether that is Amazon, Microsoft, Google, or someone else.

Regional Assets	We offer different regional computing end points for a variety of reasons. Proximity to computing assets reduces upload and download times. Our solution is designed to utilize different data centers within a designated region based on best market price at the time of the operation.
Elastic Capacity	Each operation runs on its own independent compute server(s). These compute servers are requested when the operation is started and they are terminated when the operation has been completed. There is no latency for your jobs to start. You are never queued behind other jobs. You always get a fresh environment of your own and your job starts and runs as soon as the appropriate compute server is available.
Large Genome Support	Bionano's Saphyr Compute server ships with 256 GB of memory. This is sufficient to handle nearly all <i>de novo</i> assemblies, but there are some plant / animal genomes that require more memory. The Bionano Compute On Demand solution offers compute nodes with 488 GB of memory for larger operations.
Performance	The Compute On Demand solution allocates compute resources to your job based on the volume of data to be processed. This additional compute power allows your job to process faster and cheaper.
Execution Silos	Each job runs on an independent set of compute servers. Your job is not executed on a shared resource.
No System Updates	The Bionano Compute On Demand service is a hosted service so you do not need to apply or install periodic updates. All system updates for the Bionano Compute On Demand service are managed by Bionano.
System Health	When you sign up for the Bionano Compute On Demand service your Bionano Access system will share system health metrics with Bionano not including any protected data. This will allow Bionano to better support our customer base.

Terms and Conditions

The terms and conditions that your system administrator has agreed to on your behalf in order to enable a connection to Bionano Compute On Demand and/or Saphyr Assure are below. By using Bionano Compute On Demand and/or Saphyr Assure, you are bound to these terms and conditions.

- You may only use Bionano Compute On Demand and Saphyr Assure to process data that is owned, licensed or lawfully obtained by you.
- You must comply with the current Bionano technical documentation applicable to the data you send to the Bionano Compute On Demand. Noncompliance to this technical documentation may result in failure of operations that would not be reimbursable by Bionano.
- Genomic Data (i.e. BNX files) will be transferred from your local server to a hosted Bionano Compute On Demand resource for the purpose of computing in the cloud. These resources exist in AWS and/or Microsoft Azure datacenters and not on Bionano-owned infrastructure. The data is encrypted during transit to those datacenters. Refer to 30292, Bionano Genomics Data Security Guidelines for details of the data that is transmitted.
- Upon successful completion, data are returned to your Bionano Access Server and are deleted from Bionano Compute On Demand.
- When operations fail in Bionano Compute On Demand, data are retained temporarily within the Bionano Compute On Demand servers. You may authorize Bionano Support to view this data for the sole purpose of troubleshooting the failure. Whether or not you contact Bionano, this temporary data will be deleted pursuant to Bionano procedures in effect at that time.
- Tokens provided from Bionano are required to run operations in the Bionano Compute On Demand. You will act ethically in procuring and using tokens. You will not attempt to misuse or otherwise attempt to subvert the use of tokens.
- System health metrics will be transferred from your Saphyr system to Saphyr Assure that is hosted in a Microsoft Azure datacenter. The data is encrypted during transit to this datacenter. Refer to 30292, Bionano Genomics Data Security Guidelines for details of the data that is transmitted.
- By enabling Bionano Compute On Demand and Saphyr Assure, you authorize Bionano to collect summary metrics to support billing, support and product improvement. These metrics include information on the status of your Bionano Access and Saphyr system and are not shared with third parties. Bionano does not collect any protected health information nor your actual run data or pipeline results. You can select Saphyr Assure Only without Compute On Demand. This enables the collection of the same summary metrics, but not allowing your genomic data (i.e. BNX file) to be sent to Bionano Compute On Demand.
- You agree to receive emails from Bionano pertaining to Bionano Compute On Demand and Saphyr Assure system maintenance and upgrade events. We do not share your email with any third parties. These emails will be sent to the email addresses registered in your Bionano Access server and/or associated with your account within Bionano's systems.
- Bionano, at its sole discretion and without prior notice, may terminate your access to the Bionano Compute On Demand and Saphyr Assure if you are not in compliance with the terms of use or otherwise behave in an illegal or unethical way.

What's New in this Release

- New rates have been established with our cloud providers. This should provide a moderate savings on all operations across the board.
- Human *de novo* assemblies for Access 1.6 have been moved to a new compute platform. They now execute roughly 25% faster and cost 50% less. This new compute platform is not yet available in the Germany compute region.
- The latest version of Access offers tremendous Compute On Demand savings over previous versions. As such, we have made our estimates version specific. This allows us to require fewer reserved tokens to start operations on the Compute On Demand service for newer versions.

Update Tickets

Version	Summary	Ticket
1.1	Enable 5Tbp RVP	STRAT-346
1.1	Switch to High Performance Storage Architecture	STRAT-278
1.1	Increase maximum DNA permitted for De Novo Assembly to 4Tbp	STRAT-319
1.1	Add BNX Alignment Operation	STRAT-26
1.2	Move Human De Novo to Azure	STRAT-473
1.2	New Rate System	STRAT-491
1.2	Pipeline Specific Estimates	STRAT-394
1.2.1	Incomplete Rare Variant Analysis Issue Fixed	STRAT-583

Token Ranges

The table below identifies the token ranges expected for each operation in the current release. There will be fluctuations in token pricing from release to release as we continue to reduce costs and accelerate performance.

Operation	min tokens	max tokens	Comment
Human de Novo Assembly	1	Calculated	max tokens computed based on input molecules. 2 Tbp limit (down sampled to 250X at runtime)
Non-Human / No Reference de Novo Assembly	1	50	Below 300 Gbp
Non-Human / No Reference de Novo Assembly	10	150	300 – 500 Gbp
Non-Human / No Reference de Novo Assembly	20	250	500 – 1000 Gbp
Non-Human / No Reference de Novo Assembly	30	400	1000 – 2000 Gbp
Non-Human / No Reference de Novo Assembly	40	500	2000 – 3000 Gbp (Limit increased)
Non-Human / No Reference de Novo Assembly	50	650	3000 – 4000 Gbp (New tier)
Scaffold	1	30	
Two Enzyme Hybrid Scaffold	1	30	
Variant Annotation (single)	1	3	
Variant Annotation (dual or trio)	1	10	
Bnx Merge	1	2	
Filter Bnx	1	2	
MQR	1	2	
SV Merge	1	5	
Alignment	1	4	
Rare Variant Analysis	3	25	Below 2000 Gbp
Rare Variant Analysis	5	50	2000 – 6000 Gbp (New tier)
FSDH	1	5	New operation

Technical Assistance

For technical assistance, contact Bionano Genomics Technical Support.

You can retrieve documentation on Bionano products, SDS's, certificates of analysis, frequently asked questions, and other related documents from the Support website or by request through e-mail and telephone.

Type	Contact
Email	support@bionanogenomics.com
Phone	Hours of Operation: Monday through Friday, 9:00 a.m. to 5:00 p.m., PST US: +1 (858) 888-7663
Website	www.bionanogenomics.com/support