

2015-03-27 V1.0





# **Contents**

1 Overview	4
1.1 Introduction	4
1.2 Purpose	4
1.3 Intended Audience	4
1.4 Service Scenarios	5
1.5 Customer Benefits	<i>6</i>
1.6 Key Components	7
1.7 Load Model	7
2 Products and Technologies	10
2.1 Huawei OceanStor 9000 Big Data Storage System	10
2.2 Huawei FusionServer E9000 Blade Server	11
2.3 Illumina Isaac Aligner and Variant Caller	11
3 Best Practices for Illumina Genome Analysis Based on Huawei OceanStor 9000	12
3.1 Service Challenges	
3.2 Huawei Solution	12
3.2.1 Solution Architecture	12
3.2.2 Solution Configuration	13
3.3 Verification Procedure	15
3.3.1 Overview	15
3.3.2 System Planning	15
3.3.3 System Deployment	18
3.3.4 System Verification	18
3.4 Verification Results	
3.5 Result Analysis	
3.5.1 Improving the Data Analysis Performance by 212%	23
3.5.2 Completing One Round of HiSeq X10 Whole-Genome Analysis in 2 Days and 8 Hours	24
4 Typical Configuration	<b>2</b> 5
4.1 Planning and Configuration	25
4.2 Product Constraints	26
5 Appendix	27
5.1 Appendix List	27



5.2 Reference Documents	27
5.3 Acronyms and Abbreviations	27



#### 1.1 Introduction

With the rapid development of society, economy, science, and technology as well as improvement of living standards, people pay more attention to exploration of life science and pursuit of healthy life. Technologies such as genome sequencing and genome analysis are highly promoted and widely used in fields such as health, medical care, and politics and law. However, limited performance and scale-out capabilities prevent traditional IT systems from addressing the fast-growing requirements for genome sequencing. The industry is in urgent need of new IT system solutions for storage and analysis of massive genome data.

With the assistance of bioinformatics specialists from Illumina, Huawei has built an IT system for genome sequencing based on the OceanStor 9000 big data storage system and FusionServer E9000 blade server. In addition, Huawei has tested and verified the IT system.

This document describes the best practices for Illumina genome analysis based on Huawei OceanStor 9000 big data storage system and Huawei FusionServer E9000 blade server. The key workflow and algorithm of Illumina genome analysis are implemented by using the Isaac Genome Alignment (Isaac Aligner for short) and Isaac Variant Caller:

- Isaac Aligner, a piece of genome sequence alignment software, compares human genome indexes and outputs a result file in BAM format. Such a file acts as the input for subsequent variant calling, structural variation calling, and copy number variation (CNV)
- Isaac Variant Caller, a piece of variant calling software, analyzes the results provided by the Isaac Aligner and provides analysis results in VCF format for further analysis.

#### 1.2 Purpose

This document provides Huawei partners and customers with an end-to-end IT solution for Illumina genome analysis. It aims to provide a technical reference for customers who want to deploy Huawei OceanStor 9000 big data storage system and FusionServer E9000 blade server to build an IT support system for genome analysis, helping Huawei partners and customers to simplify IT solution designs.

#### 1.3 Intended Audience

This document is intended for Huawei partners, customers, and employees, especially IT managers of customers who want to deploy Huawei OceanStor 9000 big data storage system and FusionServer E9000 blade server to build an IT support system for genome analysis.



The following knowledge can help readers better understand this document:

- Architecture and working principles of Huawei OceanStor 9000 big data storage system
- Architecture and working principles of Huawei FusionServer E9000 blade server
- Linux operating system basics
- Basics of genome sequencing and analysis

#### 1.4 Service Scenarios

The Illumina genome sequencing platform, such as the Illumina HiSeq X10 system, collects raw genome data and imports the raw data to the storage system through the high-speed switching network for further analysis.

The Illumina genome analysis system consists of a switching network, computing cluster, and storage system.

- The switching network connects the genome sequencing platform, computing cluster, and storage system, and exchanges data between these subsystems.
- The computing cluster is deployed with diversified genome analysis software to read raw genome data generated by the genome sequencing platform from the storage system through the switching network, analyze the data, and write the process data and result files back to the storage system.
- The storage system stores the raw genome data, as well as the process data and result files of genome analysis software.

Figure 1-1 Service scenario of Illumina genome sequencing and analysis

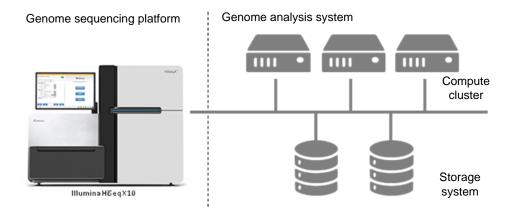
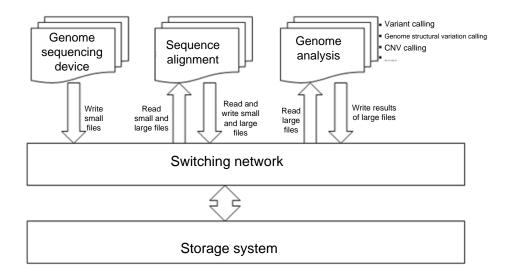




Figure 1-2 Service flow of Illumina genome sequencing and analysis



The Illumina genome sequencing and analysis service is implemented in three phases:

Phase 1: The Illumina genome sequencing device collects and processes genome data, and generates original genome files in BCL format.

Phase 2: The Illumina Isaac Aligner reads the genome files in BCL format and human reference genome indexes, and directly generates a genome description file in BAM format after sequence alignment.

Phase 3: The Illumina Isaac Variant Caller or other genome analysis software (such as the Manta genome structural variation calling software and CNV nator genome copy number variation calling software) reads the BAM file, analyzes the data, and eventually generates a result file in VCF format.

For the genome sequence alignment process in Phase 2, the traditional sequence alignment software can only analyze FASTQ files. Therefore, the format conversion software must be separately deployed to convert BCL files into FASTQ files. Unlike the traditional sequence alignment software, the Illumina Isaac Aligner does not need to convert BCL files into FASTQ files, shortening the time required for sequence alignment.

#### 1.5 Customer Benefits

The highly-integrated IT structure for genome analysis consists of the computing, storage, and switching subsystems that feature high density, high performance, and linear expansion. It can well address the requirements for storing, analyzing, and archiving massive genome data, and its performance is improved more than twice that of the standard computing system recommended in the Illumina whitepaper (issued on July 2013). For details, see Table 1-1. In addition, customers are offered the best after-sales service and assurance, which significantly reduces the maintenance costs of the system.



**Table 1-1** Comparison of the genome analysis performance

Genome Analysis System			J		Performance Improvement Achieved by Huawei IT Architecture for Genome Analysis
IlluminaCompute standard system	30x human genome data: 90 GB	90 GB/7 hours 12 minutes = 12.5 GB/hour	212%		
IlluminaCompute optimized system 30x human genome data: 90 GB		90 GB/3 hours 40 minutes = 24.55 GB/hour	59%		
Huawei IT architecture for genome analysis  Single-channel human genome data: 130 GB		130 GB/3 hours 20 minutes = 39 GB/hour	-		

Based on the Huawei IT architecture for genome analysis, a 10-node computing cluster is deployed to execute tasks concurrently. It takes only 56 hours (2 days and 8 hours) for the HiSeq X10 to complete a round of genome data analysis in 8 x 2 x 10 = 160 channels (assume that each channel contains 130 GB human genome data).

#### M NOTE

- [1] The detailed configuration of the IlluminaCompute standard system is as follows: AMD Opteron<sup>TM</sup> Processor 6212 32CPU, 128 GB, and local RAID6.
- [2]The detailed configuration of the IlluminaCompute optimized system is as follows: Intel® Xeon® CPU E5-2687 32CPU, 128 GB, and local SSDs.
- [3] The detailed configuration of the Huawei IT architecture for genome analysis is as follows: Huawei FusionServer E9000 blade server (Intel® Xeon® CPU E5-2680 32CPU, 128 GB) + Huawei OceanStor 9000 big data storage system (6\*P36 nodes).

## 1.6 Key Components

Key components include:

- Huawei OceanStor 9000 big data storage system
- Huawei FusionServer E9000 blade server
- CentOS 5.11 x86\_64 operating system
- Illumina Isaac Aligner and Isaac Variant Caller genome analysis software

#### 1.7 Load Model

A total of 80 GB raw genome data is generated in a single channel of the Illumina HiSeq X10. After the Isaac Aligner is used for sequence alignment, temporary DAT files and BAM result files with a total size of 384 GB are generated. Then, the Isaac Variant Caller implements variant calling and generates **vcf.gz**, which is a 2.5 GB result file.



For the bottom-layer storage OceanStor 9000, the entire process is a big file sequential read/write process. The read/write ratio of the Isaac Aligner is 1:1, and the read/write ratio of the Isaac Variant Caller is 40:1. Table 1-2 summarizes the entire service model.

Table 1-2 Service model of Illumina genome analysis

Item	Content
Service characteristics	Illumina Isaac Aligner 01.14.11.11: It supports sequential read/write operations. The read/write ratio is 1:1. It reads the raw data and human genome indexes, generates a lot of temporary DAT files (size of a single DAT file: 200 MB to 15 GB), and finally combines the DAT files into a super-large BAM result file (the size of the BAM result file depends on the size of the raw data and is generally similar to that of the raw data).  Illumina Isaac Variant Caller 1.0.6: It supports sequential read/write operations. The read/write ratio is 40:1. It reads the BAM file and genome.fa generated by the Isaac Aligner, and generates vcf.gz, which is a 2.5 GB result file.  If both Isaac Aligner and Isaac Variant Caller are deployed on a single client, the genome analysis tasks are executed sequentially in a serial manner. The tasks are concurrently executed on ten clients.
Total amount of data	Human genome indexes: 50 GB A single set of HiSeq X10: 466.5 GB x 160 = 72.89 TB (one analysis)
File size	<ol> <li>Human genome indexes: 50 GB, including one genome.fa file (3 GB) and 64 DAT files. The size of each DAT file ranges from 100 MB to 1.7 GB.</li> <li>Amount of data in a single channel of the HiSeq X10</li> <li>Raw genome data: 80 GB, including 29760 bcl.gz compressed files and 96 filter files. The size of each bcl.gz compressed file ranges from 2 MB to 3 MB. The size of each filter file is about 6 MB.</li> <li>Illumina Isaac Aligner 01.14.11.11: 384 GB, including 300 GB temporary DAT files and one 84 GB BAM result file. The size of each temporary DAT file ranges from 200 MB to 15 GB.</li> <li>Illumina Isaac Variant Caller 1.0.6: 2.5 GB, which is the size of a vcf.gz result file.</li> </ol>
I/O size	Size of the I/O block on the client: configurable. The recommended size is 1 MB.  Size of the application block: not determined yet.
Trend of the data amount	A total of 72.89 TB data is generated in each sequencing period, including:  1) 46.875 TB temporary files, which are deleted as soon as sequencing is completed.  2) 13.52 TB result files, which are permanently stored.  3) 12.495 TB raw data, which is stored for two weeks or longer if the business model is used.



Item	Content	
Throughout	Concurrent peak read throughput of 10 clients: Concurrent peak write throughput of 10 clients:	
Throughput (MB/s)	Note: The performance requirements vary with Here, the configuration of a client is as follows: E5-2680 32CPU, 128GB.	
	Illumina Isaac Aligner 01.14.11.11	
	Reading of the raw genome data (such as bcl, fil files) 64	ter, and <b>genome.fa</b>
Concurrency	Genome sequencing	Number of OS cores
requirement	Writing of temporary files (matches.dat and bin	<b>n.dat</b> ) 64
(Concurrent	Reading of temporary files (matches.dat and bi	in.dat) 8
clients/threads of	Output of result files (sorted.bam)	8
upper-layer applications)	Illumina Isaac Variant Caller 1.0.6	
applications)	Reading of genome data ( <b>sorted.bam</b> and <b>geno</b> ) Number of OS cores	me.fa)
	Output of result files (vcf.gz) cores	Number of OS
Latency requirement	NA	
Reliability requirement	4 + 2	
NAS service protocol type	Genome analyzer: Windows 7 Pro SMB2.1 Genome analysis software: user configuration, N	NFS V3 or V4



# **Products and Technologies**

# 2.1 Huawei OceanStor 9000 Big Data Storage System

Huawei OceanStor 9000 big data storage system uses a symmetric distributed architecture to deliver cutting-edge performance, scale-out capabilities, and super-large single file systems, providing shared storage for structured data and unstructured data. The OceanStor 9000 is applicable to various big data scenarios, including broadcasting & media, satellite mapping, genome sequencing, energy exploration, scientific research, and education.

The OceanStor 9000 uses the OceanStor distributed file system (DFS) to combine all physical storage device nodes (nodes for short) into a cluster to provide such functions as file access and capacity management. The OceanStor 9000 provides three types of nodes.

For more information about this product, click the following hyperlink:

http://e.huawei.com/cn/products/cloud-computing-dc/storage/massive-storage/9000



Figure 2-1 Huawei OceanStor 9000 big data storage system



#### 2.2 Huawei FusionServer E9000 Blade Server

Huawei FusionServer E9000 blade server is a next-generation powerful infrastructure platform that integrates computing, storage, switching, and management subsystems. The E9000 is an industry-leading hardware computing platform that improves competitiveness based on its availability, computing density, energy consumption, low emissions, backplane bandwidth, intelligent management and services, elastic configuration, flexible computing and storage expansion capabilities, low network latency, and acceleration functions.

For more information about this product, click the following hyperlink:

http://e.huawei.com/cn/products/cloud-computing-dc/servers/e-series/e9000-sever-chassis





## 2.3 Illumina Isaac Aligner and Variant Caller

The Isaac genome analysis software introduced by Illumina provides an efficient and accurate resequencing workflow. The workflow begins with BCL or FASTQ files and produces BAM or VCF files in just over 7 hours for a 30x human genome on a standard computing server (see the IlluminaCompute standard system specified in Table 1-1). This workflow provides a four- to seven-fold speed increase over currently available genome analysis software, using PCR-free sample prepared by the HiSeq 2500 system.

The Isaac Aligner is targeted at DNA sequencing with low error rates and read lengths (single or paired ends), which is typical for genome data produced by the HiSeq platform. It uses an indexing scheme to quickly perform retrieval. The Isaac Aligner runs on a single compute node and efficiently uses the CPU power while minimizing the I/O.

The Isaac Variant Caller identifies the genotypes of single-nucleotide variants (SNVs) and small variations in the case of diploid genome. The output VCF file captures the genotype at each position. These genotypes may be the genes that are different from the human genome indexes.



# **Best Practices for Illumina Genome** Analysis Based on Huawei OceanStor 9000

# 3.1 Service Challenges

The Illumina genome analysis service raises rigid requirements for the storage performance, storage capacity, and computing capability of the IT system:

- A large amount of data to be stored: A set of HiSeq X10 completes one round of sequencing in three days, during which 72.89 TB data is generated. Among the 72.89 TB data, 26 TB data needs to be stored permanently or for a long term. The data generated in a year requires a usable storage space of 3 PB.
- A large amount of data to be computed: A total of 20 TB base data is generated in one round of sequencing in three days. Only a high-performance computing cluster can complete processing of these data.
- High throughput: The genome analysis process requires a large storage capacity and a high network bandwidth. The storage system must have a minimum throughput of 30 GB/s. In addition, genome analysis requires a low latency; otherwise, the data integrity cannot be ensured.

#### 3.2 Huawei Solution

#### 3.2.1 Solution Architecture

The genome analysis system uses Huawei OceanStor 9000 big data storage system as the basic storage platform for storing all the genome data, and uses Huawei FusionServer E9000 blade server to build a high-density and high-performance computing cluster. In addition, an internal 10GE high-speed switching system is built. The storage system directly accesses the switching system for local high-speed switching. The management network uses Huawei S5700 campus switch to access the customer's maintenance network.

The genome sequencing platform also uses Huawei S5700 campus switch to access the genome analysis network through the 20GE trunk in the uplink.



Figure 3-1 Logical architecture of the solution

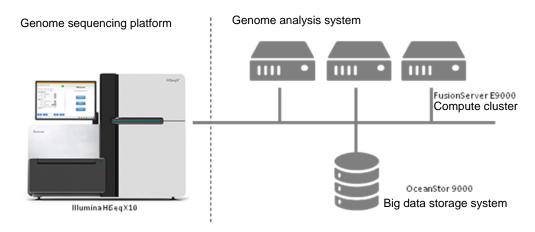
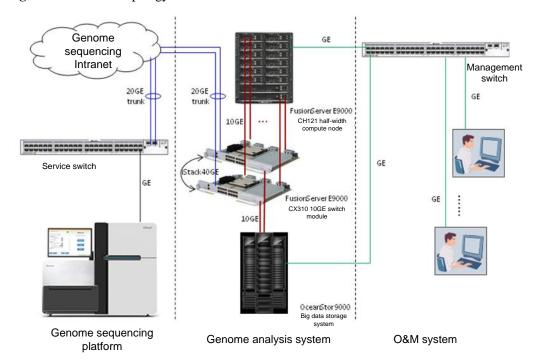


Figure 3-2 Network topology



#### **Ⅲ** NOTE

- [1] The uplink switching capacity can be expanded based on the service requirements.
- [2] Figure 3-2 shows the complete network topology of the solution, but the verification does not cover the genome sequencing platform.

#### 3.2.2 Solution Configuration

For the computing cluster, the E9000 CH121 half-width compute nodes are configured, and up to 16 nodes can be configured for a single chassis. For the switching system, the E9000-embedded CX310 10GE high-speed switch module is configured. For the storage



system, the P36 node of the OceanStor 9000 is configured, and a single node can be fully configured with 35 enterprise-level SATA disks and one SSD, with a raw capacity of 140 TB.

Table 3-1 describes the hardware configuration, and Table 3-2 describes the software configuration.

 Table 3-1 Hardware configuration

Device Type	Hardware Type	re Type Component Specifications	
OceanStor 9000 big data storage system	P36 node	35 x 4 TB enterprise-level SATA disks + 1 x 200 GB SSD	6
FusionServer E9000 blade	CH121 half-width compute node	2 x Intel® Xeon® CPU E5-2680 0 @ 2.70 GHz + 8 x DDR3 RDIMM-16 GB-1600 MHz + 1 x Emulex OneConnect 10Gb*2 NIC	10
server	CX310 10GE converged switch module	16 x uplink 10GE SFP+ ports + 32 downlink 10GE ports	2
Huawei campus switch	S5700-52C-EI	2 x uplink 10GE SFP+ ports + 48 downlink 10/100/1000Base-T ports	1

 Table 3-2 Software configuration

Device Type	Hardware Type	Software
OceanStor 9000 big data storage system	P36 node	OceanStor 9000 V100R001
FusionServer E9000 blade server	CH121 half-width compute node  CX310 10GE converged switch module	CentOS 5.11 x86_64 operating system  Parameter optimization: echo 16777216 16777216 16777216 > /proc/sys/net/ipv4/tcp_mem echo 16777216 16777216 16777216 > /proc/sys/net/ipv4/tcp_rmem echo 16777216 16777216 16777216 > /proc/sys/net/ipv4/tcp_wmem sysctl -w vm.dirty_bytes=167772160  Illumina Isaac Aligner 01.14.11.11  Illumina Isaac Variant Caller 1.0.6
Huawei campus switch	S5700-52C-EI	S5700 V100R005

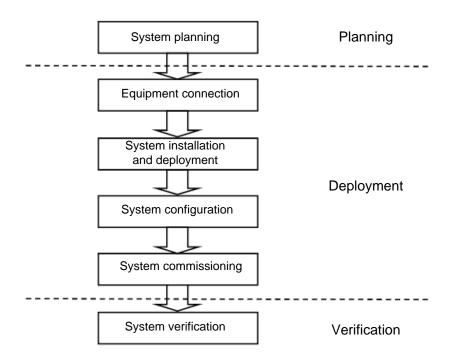


#### 3.3 Verification Procedure

#### 3.3.1 Overview

The entire genome analysis verification process includes planning, deployment, and verification, as shown in Figure 3-3.

Figure 3-3 Genome analysis verification procedure



#### 3.3.2 System Planning

#### 3.3.2.1 Network Planning

The Huawei IT architecture for genome analysis consists of the service network (read/write access to the genome analysis data), storage network (internal switching of the storage system), and management system (service operations and equipment maintenance), as described in Table 3-3.

Table 3-3 Network planning of the genome analysis system

Device Type	Network Type	Network Segment	Gateway	IP Address
OceanStor	Storage	129.203.10		129.203.103.11 to
9000 big	network	3.0/24		129.203.103.22



Device Type	Network Type	Network Segment	Gateway	IP Address
data storage system	Service network	129.28.103. 0/24	129.28.103.2 54	InfoEqualizer DNS: 129.28.103.100 Service-specific static IP address: 129.28.103.11 to 129.28.103.22 Service-specific dynamic IP address: 129.28.103.51 to 129.28.103.62
	Management network	100.28.103. 0/24	100.28.103.2 54	IP address for system management: 100.28.103.100 IP address for equipment IPMI management: 100.28.103.1 to 100.28.103.6
	Service network	129.28.103. 0/24	129.28.103.2 54	Set the balance-xor bond mode for two ports of each node: 129.28.103.101 to 129.28.103.110
FusionServe r E9000 blade server	Management network	100.28.103. 0/24	100.28.103.2 54	IP address for BMC management: 100.28.103.101 to 100.28.103.110 IP address for SMM management: 100.28.103.10 to 100.28.103.12 IP address for switch board management: 100.28.103.13 to 100.28.103.14
Huawei campus switch	Management network	100.1.201.0 /24		Management IP address: 100.1.201.1

#### 3.3.2.2 Storage Planning

A set of Illumina HiSeq X10 consists of ten sequencing devices. Each sequencing device has two chips. The data of each chip is output to a storage system directory through the CIFS.



These directories are named XtenAA, XtenAB, XtenBA, XtenBB, and so on. A total of 20 directories are needed. The human reference genome data is separately stored in the **iGenomes** directory, so that all the compute nodes can access the data.

According to the planning, ten compute nodes are configured, and mapped to sequencing devices on a one-to-one basis. These nodes can access the directory of the raw genome data on the storage system through the NFS interface. The load is manually and evenly distributed to each node of the storage system based on the service model and traffic.

**Table 3-4** Storage planning

Table 5-4 Storage planning				
Sequencing Directory on the Storage System		Mount Command		
XtenA	XtenAA XtenAB	mount -t nfs -o noatime,nodiratime,noacl,rw,vers=3,wsize=1048576,rsi ze=1048576 129.28.103.56:/xtenAA /mnt/xtenaa mount -t nfs -o noatime,nodiratime,noacl,rw,vers=3,wsize=1048576,rsi ze=1048576 129.28.103.62:/xtenAB /mnt/xtenab mount -t nfs -o noatime,nodiratime,noacl,rw,vers=3,wsize=1048576,rsi ze=1048576 129.28.103.52:/iGenomes /mnt/iGenomes		
XtenB	XtenBA XtenBB	mount -t nfs -o noatime,nodiratime,noacl,rw,vers=3,wsize=1048576,rsi ze=1048576 129.28.103.57:/xtenBA /mnt/xtenba mount -t nfs -o noatime,nodiratime,noacl,rw,vers=3,wsize=1048576,rsi ze=1048576 129.28.103.51:/xtenBB /mnt/xtenbb mount -t nfs -o noatime,nodiratime,noacl,rw,vers=3,wsize=1048576,rsi ze=1048576 129.28.103.58:/iGenomes /mnt/iGenomes		
XtenI	XtenIA XtenIB	mount -t nfs -o noatime,nodiratime,noacl,rw,vers=3,wsize=1048576,rsi ze=1048576 129.28.103.55:/xtenIA /mnt/xtenia mount -t nfs -o noatime,nodiratime,noacl,rw,vers=3,wsize=1048576,rsi ze=1048576 129.28.103.61:/xtenIB /mnt/xtenib mount -t nfs -o noatime,nodiratime,noacl,rw,vers=3,wsize=1048576,rsi ze=1048576 129.28.103.52:/iGenomes /mnt/iGenomes		
XtenJ	XtenJA XtenJB	mount -t nfs -o noatime,nodiratime,noacl,rw,vers=3,wsize=1048576,rsi ze=1048576 129.28.103.54:/xtenJA /mnt/xtenja mount -t nfs -o noatime,nodiratime,noacl,rw,vers=3,wsize=1048576,rsi ze=1048576 129.28.103.60:/xtenJB /mnt/xtenjb mount -t nfs -o		



Sequencing Dovice No.	Corresponding Directory on the Storage System	Mount Command
		noatime,nodiratime,noacl,rw,vers=3,wsize=1048576,rsi ze=1048576 129.28.103.58:/iGenomes /mnt/iGenomes

#### MOTE

[1] The data protection level of the OceanStor 9000 is N + 2, the strip size is 128 KB, and the data can be read or written by the genome analysis software and sequencing devices through the NFS v3 and CIFS (SMB 2.1).

#### 3.3.3 System Deployment

- Connect the network based on the network topology and planning of the system, and configure the switching network based on the Huawei FusionServer CX310 Switch Module Configuration Guide and S5700 Switch Configuration Guide.
- Install and configure the storage system based on the Huawei OceanStor 9000 Hardware Installation Guide, Huawei OceanStor 9000 Software Installation Guide, and system planning.
- Install the computing cluster, and install and configure the operating system of the compute nodes based on the FusionServer E9000 Installation Guide and system planning.
- Compile, deploy, and configure the genome analysis software on the compute nodes based on the related compilation and deployment guides of Illumina Isaac software and system planning.
- Perform system commissioning to ensure that the genome analysis software can run normally.
- Input the raw genome data and human reference genome data to the storage system, and start the test.

#### 3.3.4 System Verification

#### Item 1: Verifying the Time Required by a Compute Node to Analyze Raw Genome Data in a Single Channel

Step 1 Start the Isaac Aligner for sequence alignment. Record the start time, the completion time, and the bandwidth and CPU usage of the storage system.

Command line: /opt/isaac\_aligner/bin/isaac-align -r /mnt/iGenomes/Homo sapiens/UCSC/hg19/Sequence/IsaacIndex/sorted-reference.xml -m 120 -b /mnt/xtenaa/RunInfo.xml --base-calls-format bcl-gz -o /mnt/xtenaa/isaac-align -t /mnt/xtenaa/isaac-align/Temp



Bandwidth 570,060 456.048 342,036 228,024 114,012 0 11:06:00 11:16:00 11:46:00 11:56:00 Bandwidth (KB/s) Write bandwidth (KB/s) Read bandwidth (KB/s)

Figure 3-4 System bandwidth when the Isaac Aligner is started

Figure 3-5 System load when the Isaac Aligner is started

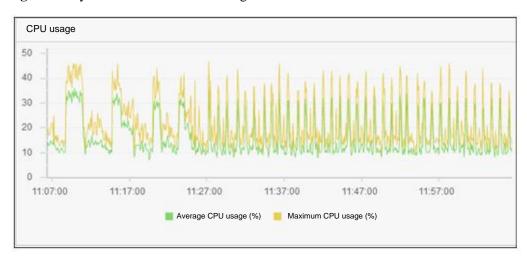


Figure 3-4 and Figure 3-5 show the bandwidth and CPU usage of the storage system when the Isaac Aligner is started. When the Isaac Aligner is started, it reads a lot of raw genome data and human reference genome data, and generates temporary files. Therefore, the performance requirement of the storage system is the highest. Later on, the Isaac Aligner periodically performs sequence alignment and periodically reads or writes the storage system.



Bandwidth 553,660 442,928 332,196 221.464 110,732 12:58:00 13:08:00 13:18:00 13:28:00 13:38:00 13:48:00 Bandwidth (KB/s) Write bandwidth (KB/s) Read bandwidth (KB/s)

Figure 3-6 System bandwidth when the Isaac Aligner is stopped

Figure 3-7 System load when the Isaac Aligner is stopped

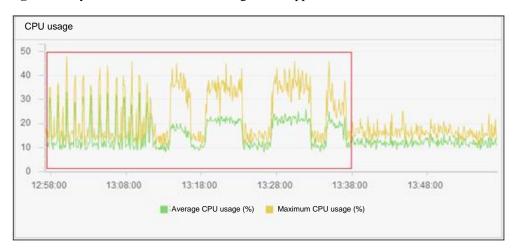


Figure 3-6 and Figure 3-7 show the bandwidth and CPU usage of the storage system when the Isaac Aligner is stopped. When the Isaac Aligner is stopped, it needs to read a lot of temporary files and generate a final result file. Therefore, the performance requirement of the storage system is also very high.

**Step 2** Prepare the software running configuration and parameters of the Isaac Variant Caller.

Command line: mkdir /mnt/xtenaa/isaac\_variant\_caller

cp/opt/isaac\_variant\_caller/etc/ivc\_config\_default.ini /mnt/xtenaa/isaac\_variant\_caller/config.ini

/opt/isaac\_variant\_caller/bin/configureWorkflow.pl

- --bam=/mnt/xtenaa/isaac-align/Projects/default/default/sorted.bam
- --ref=/mnt/iGenomes/Homo\_sapiens/UCSC/hg19/Sequence/WholeGenomeFasta/genome.fa
- --config=/mnt/xtenaa/isaac\_variant\_caller/config.ini
- --output-dir=/mnt/xtenaa/isaac\_variant\_caller/myAnalysis

Step 3 Start the Isaac Variant Caller for variant calling. Record the start time, the completion time, and the bandwidth and CPU usage of the storage system.



Command line: make -C /mnt/xtenaa/isaac\_variant\_caller/myAnalysis -j 32

Figure 3-8 System bandwidth when the Isaac Variant Caller is running

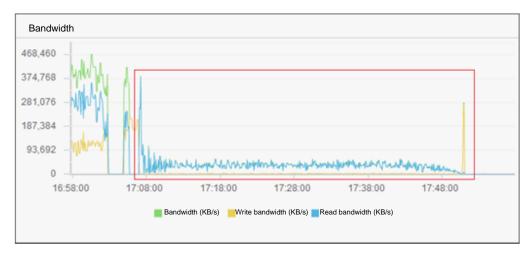


Figure 3-9 CPU usage when the Isaac Variant Caller is running

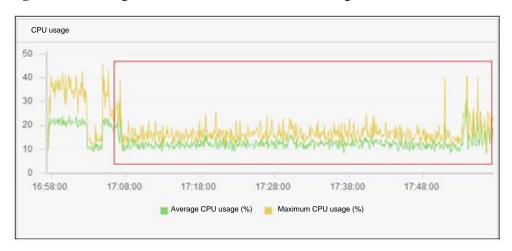


Figure 3-8 and Figure 3-9 show the bandwidth and CPU usage of the storage system when the Isaac Variant Caller is running. Compared with the Isaac Aligner, the Isaac Variant Caller only reads data when started and writes a result file when stopped, and the operation duration is very short. Therefore, the performance requirement of the storage system is not high.

#### Item 2: Verifying the Time Required by Ten Compute Nodes to Analyze Raw Genome Data in 160 Channels

Start the Isaac Aligner and Isaac Variant Caller as instructed in Item 1 on ten nodes at the same time for parallel analysis. For each node, analyze the genome data of  $8 \times 2 = 16$ channels of each sequencing device in a serial manner. Record the start time, the completion time, and the bandwidth and CPU usage of the storage system.



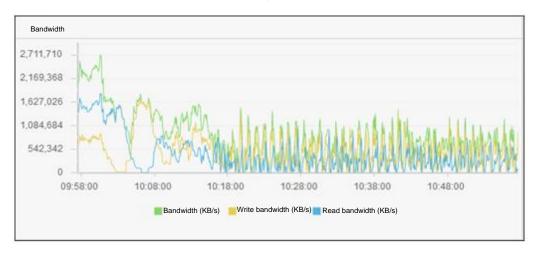
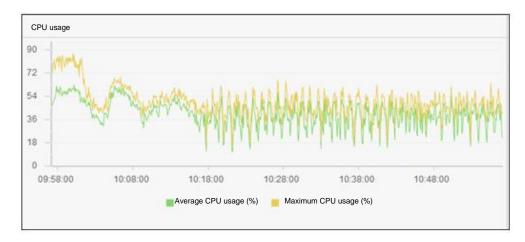


Figure 3-10 System bandwidth when the Isaac genome analysis software is started

Figure 3-11 CPU usage when the Isaac genome analysis software is started



#### 3.4 Verification Results

To prevent inaccurate test results due to variations in the system running environment and status and to ensure reliability of the test results, the testing personnel has tested the preceding two items for multiple times, and used the average values as the final test results.

#### Item 1: A Single Compute Node Completes Sequencing of Genome Data in a **Single Channel**

This item is tested three times. Each channel contains 130 GB genome data. Table 3-5 lists the test results.



**Table 3-5** Run time records of a single node

SN.	Phase Point	Time	Duration	Isaac Run Time
	isaac_aligner start time	15:28:12		
	isaac_aligner end time	18:00:59	2:32:47	
1	isaac_variant_caller start time	18:01:38		
	isaac_variant_caller end time	18:48:01	0:46:23	3:19:10
	isaac_aligner start time	18:50:59		
	isaac_aligner end time	21:24:15	2:33:16	
2	isaac_variant_caller start time	21:24:31		
	isaac_variant_caller end time	22:11:02	0:46:31	3:19:47
	isaac_aligner start time	22:20:48		
	isaac_aligner end time	0:54:26	2:33:38	
3	isaac_variant_caller start time	0:54:42		
	isaac_variant_caller end time	1:40:53	0:46:11	3:19:49

#### Item2: Ten Compute Nodes Complete Sequencing of Genome Data in 160 Channels

This item is tested three times. Each channel contains 130 GB genome data. Table 3-6 lists the test results.

Table 3-6 Run time records of ten nodes that run concurrently

SN.	Isaac Run Time
1	2 days, 7 hours, and 55 minutes
2	2 days, 8 hours, and 12 minutes
3	2 days, 8 hours, and 5 minutes

# 3.5 Result Analysis

## 3.5.1 Improving the Data Analysis Performance by 212%

Verification results indicate that it takes about 3 hours and 20 minutes for the Huawei IT architecture for genome analysis and Illumina Isaac genome analysis software to analyze 130 GB genome data in a single channel. Compared with the standard computing system



recommended in the Illumina white paper, the Huawei IT architecture for genome analysis improves the genome data analysis performance by 212%. Figure 3-4 also indicates the peak read/write bandwidth of the storage system required for the single-node genome analysis is about 557 MB/s under the current configuration.

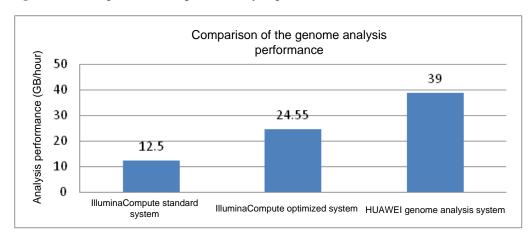


Figure 3-12 Comparison of the genome analysis performance

#### M NOTE

- [1] Huawei FusionServer E9000 blade server can be configured with higher-performance compute
- [2] For servers with the same configuration, it takes about 7 minutes and 30 seconds for the bcl2fastq format conversion software to convert a BCL file into a FASTQ file.

## 3.5.2 Completing One Round of HiSeq X10 Whole-Genome Analysis in 2 Days and 8 Hours

Based on the Huawei IT architecture for genome analysis and Illumina Isaac genome analysis software, it takes about 2 days and 8 hours for ten compute nodes to concurrently analyze genome data generated in 8 x 2 x 10 = 160 channels of the Illumina HiSeq X10. On average, it takes about 3 hours and 30 minutes to analyze 130 GB genome data in a single channel. Therefore, the requirement for completing one round of genome sequencing by the HiSeq X10 in three days can be met.

When the ten compute nodes under the current configuration concurrently perform genome analysis, the peak read/write bandwidth of the storage system reaches 2648 MB/s (88% of the bandwidth capability of the storage system, as shown in Figure 3-10), and the peak CPU usage also reaches 88%, achieving the optimal configuration of compute nodes and storage system.



# **Typical Configuration**

# 4.1 Planning and Configuration

Table 4-1 lists the recommended optimal configuration that enables a set of Illumina HiSeq X10 to complete one round of genome analysis in three days.

Table 4-1 Optimal configuration of the Huawei IT architecture for genome analysis

Device Type	Hardware Type	Component Specifications	Quantity
OceanStor 9000 big data storage system	P36 node	35 x 4 TB enterprise-level SATA disks + 1 x 200 GB SSD	6
FusionServer E9000 blade	CH121 half-width compute node	2 x Intel® Xeon® CPU E5-2680 0 @ 2.70 GHz + 8 x DDR3 RDIMM-16 GB-1600 MHz + 1 x Emulex OneConnect 10Gb*2 NIC	10
server	CX310 10GE converged switch module	16 x uplink 10GE SFP+ ports + 32 downlink 10GE ports	2
Huawei campus switch	S5700-52C-EI	2 x uplink 10GE SFP+ ports + 48 downlink 10/100/1000Base-T ports	1

#### MOTE

• [1] Under the optimal configuration, the FusionServer CH121 half-width compute node is configured with the mid-range CPUs. CPUs with better performance can be configured based on customers' requirements. If CPUs with better performance or more compute nodes are configured, it must be evaluated whether the performance of six P36 nodes of the OceanStor 9000 can meet requirements.

The storage nodes of Huawei OceanStor 9000 can be expanded based on the service performance and capacity requirements. A single file system can support up to 60 PB data. Under the typical configuration, the available capacity of six P36 nodes is about 500 TB. If 26 TB genome data that should be stored permanently or for a long term is generated in each round of sequencing, the related configuration allows a set of HiSeq X10 to run continuously for 50 days. If the raw genome data (12.495 TB raw genome data is generated in each round



of sequencing) is stored for only two weeks, the related configuration allows a set of HiSeq X10 to run continuously for 90 days.

Huawei FusionServer E9000 blade server can be configured with up to 16 CH121 compute nodes. The compute nodes and chassis can be expanded based on service requirements, or the CH121 V3 compute node with higher performance can be used instead. In addition, the CX312 10GE converged switch module (including 24 10GE uplink ports and 32 10GE downlink ports) can also be used to meet the subsequent expansion requirement of the embedded high-speed switching system.

#### **4.2 Product Constraints**

None.



# 5 Appendix

# 5.1 Appendix List

None.

# **5.2 Reference Documents**

OceanStor 9000 Big Data Storage System Technical White Paper

Huawei E9000 Server V100R001 White Paper

Illumina, Inc. whitepaper\_isaac\_workflow

# 5.3 Acronyms and Abbreviations

Table 5-1 Acronyms and abbreviations

Acronym and Abbreviation	Full Name
SPEC	Standard Performance Evaluation Corporation. It is a third-party application performance testing organization that aims to establish, maintain and endorse standards for evaluating performance of server applications.
NFS	Network File Server
Hadoop	Hadoop is an Apache distributed system framework. Users can develop distributed programs and fully utilize the high-speed computing and storage capabilities of the cluster without knowing details about the distributed bottom layer. Hadoop can implement distributed processing of a large amount of data in a reliable, efficient, and scalable manner. Hadoop is reliable because it is assumed that the computing elements and storage will fail, and therefore maintains multiple working data copies to ensure a second distributed processing on the failure node. Hadoop is efficient because it runs in a parallel manner to increase the processing speed. Hadoop is scalable and can process data at the PB level. Hadoop



Acronym and Abbreviation	Full Name
	consists of the Hadoop distributed file system (HDFS), MapReduce, HBase, and ZooKeeper.



#### Copyright © Huawei Technologies Co., Ltd. 2015. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

#### **Trademarks and Permissions**



HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

#### **Notice**

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

#### Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base

Bantian, Longgang Shenzhen 518129

People's Republic of China

Website: http://www.huawei.com