Overview of instance families

This topic describes all Elastic Compute Service (ECS) instance families available for purchase and introduces their features, instance types, and supported scenarios to facilitate instance type selection.

Background information

Before you read further in this topic, you must be familiar with the following information:

- **Classification and naming of instance types**. Familiarize yourself with the instance family categories, naming conventions of instance types, and differences between instance families. For more information, see Classification and naming of instance types.
- **Instance type metrics**. For information about the metrics of instance types, see Instance type metrics. You can also call the DescribeInstanceTypeFamilies and DescribeInstanceTypes operations to query the instance families and the details of all instance types provided by ECS.
- **Instructions for selecting instance types based on your business scenarios** . For more information, see Instance type selection.

After you determine an instance type for your use case, you may need to learn about the following information:

- **Regions in which the instance type is available for purchase** . Instance types that are available for purchase vary based on the region. You can go to the Instance Types Available for Each Region page to view the instance types available for purchase in each region. Alternatively, you can call the DescribeRegions and DescribeZones operations to query the available regions and the zones in a specific region.
- **Estimated instance costs**. You can calculate the price of instances that uses different billing methods in the Price Calculator. You can also call the DescribePrice operation to query information about the most recent prices of ECS resources.
- **Instructions for purchasing an instance**. You can go to the ECS instance buy page to place a purchase order for instances.

You may be concerned about the following information:

- **Retired instance families**. If you cannot find an instance type in this topic, the instance type may be in a retired instance family. For information about retired instance families, see Retired instance families.
- **Supported instance type changes**. Before you change the instance type of an instance, check whether the instance type can be changed and identify compatible instance types. For more information, see Instance types and families that support instance type changes.

Catalog

x86-based enterprise-level computing instance families

General-purpose instance families (g series)

Intel processor-powered instance families	AMD processor-powered instance families	Not recommended instance families (If the following instance families are sold out, we recommend that you use the instance families in the preceding columns.)

• g8i, general-purpose instance family • g7, general-purpose instance family • g6e, performance-enhanced general-purpose instance family • g6, general-purpose instance family	• g8a, general-purpose instance family • g8ae, performance-enhanced general-purpose instance family • g7a, general-purpose instance family • g6a, general-purpose instance family	• g5, general-purpose instance family • sn2ne, network-enhanced general-purpose instance family

Compute-optimized instance families (c series)

Intel processor-powered instance families	AMD processor-powered instance families	Not recommended instance families (If the following instance families are sold out, we recommend that you use the instance families in the preceding columns.)
• c8i, compute-optimized instance family • c7, compute-optimized instance family • c6e, performance-enhanced compute-optimized instance family • c6, compute-optimized instance family	• c8a, compute-optimized instance family • c8ae, performance-enhanced compute-optimized instance family • c7a, compute-optimized instance family • c6a, compute-optimized instance family	• ic5, compute-intensive instance family • c5, compute-optimized instance family • sn1ne, network-enhanced compute-optimized instance family

Memory-optimized instance families (r series)

Intel processor-powered instance families	AMD processor-powered instance families	Not recommended instance families (If the following instance families are sold out, we recommend that you use the instance families in the preceding columns.)
• r8i, memory-optimized instance family • r7, memory-optimized instance family • r6e, performance-enhanced memory-optimized instance family • r6, memory-optimized instance family	• r8a, memory-optimized instance family • r8ae, performance-enhanced memory-optimized instance family • r7a, memory-optimized instance family • r6a, memory-optimized instance family	• r5, memory-optimized instance family • se1ne, network-enhanced memory-optimized instance family • se1, memory-optimized instance family

Universal instance families

u1, universal instance family

Big data instance families (d series)

Recommended instance families	Not recommended instance families (If the following instance families are sold out, we recommend that you use the recommended instance families.)
• d3s, storage-intensive big data instance family • d3c, compute-intensive big data instance family • d2c, compute-intensive big data instance family • d2s, storage-intensive big data instance family • d1ne, network-enhanced big data instance family	d1, big data instance family

Instance families with local SSDs (i series)

Recommended instance families			Not recommended instance families (If the following instance families are sold out, we recommend that you use the instance families in the preceding columns.)
Instance families powered by **Intel® Xeon® Scalable (Ice Lake) processors**	Instance families powered by **Intel® Xeon® Platinum 8269CY (Cascade Lake) processors**	Instance families powered by **Intel® Xeon® Platinum 8163 (Skylake) processors**	
• i4, instance family with local SSDs • i4g, instance family with local SSDs • i4r, instance family with local SSDs • i4p, performance-enhanced instance family with local SSDs	• i3g, instance family with local SSDs • i3, instance family with local SSDs	• i2, instance family with local SSDs • i2g, instance family with local SSDs • i2ne, instance family with local SSDs • i2gne, instance family with local SSDs	i1, instance family with local SSDs

Instance families with high clock speeds (hf series)

Recommended instance families		Not recommended instance families (If the following instance families are sold out, we recommend that you use the instance families in the preceding columns.)
Instance families powered by **Intel® Xeon® Cooper Lake processors**	Instance families powered by **Intel® Xeon® Platinum 8269CY (Cascade Lake) processors**	
• hfc7, compute-optimized instance family with high clock speeds • hfg7, general-purpose instance family with high clock speeds • hfr7, memory-optimized instance family with high clock speeds	• hfc6, compute-optimized instance family with high clock speeds • hfg6, general-purpose instance family with high clock speeds • hfr6, memory-optimized instance family with high clock speeds	• hfc5, compute-optimized instance family with high clock speeds • hfg5, general-purpose instance family with high clock speeds

Enhanced instance families

Storage-enhanced instance families	Network-enhanced instance families	Security-enhanced instance families	Memory-enhanced instance families

• g7se, storage-enhanced general-purpose instance family
• c7se, storage-enhanced compute-optimized instance family
• r7se, storage-enhanced memory-optimized instance family
	• g7nex, network-enhanced general-purpose instance family
	• c7nex, network-enhanced compute-optimized instance family
	• g7ne, network-enhanced general-purpose instance family
	• g5ne, network-enhanced general-purpose instance family
		• g7t, security-enhanced general-purpose instance family
		• c7t, security-enhanced compute-optimized instance family
		• r7t, security-enhanced memory-optimized instance family
		• g6t, security-enhanced general-purpose instance family
		• c6t, security-enhanced compute-optimized instance family
			• **Recommended instance families**
			  ◦ re6p, persistent memory-optimized instance family
			  ◦ re6, high-memory instance family
			• **Not recommended instance families (If the following instance families are sold out, we recommend that you use the recommended instance families.)**
			  ◦ re4, high-memory instance family
			  ◦ re4e, high-memory instance family

x86-based entry-level computing instance families

Recommended instance families	Not recommended instance families (If the following instance families are sold out, we recommend that you use the recommended instance families.)
• e, economy instance family • t6, burstable instance family	• t5, burstable instance family • v5, CPU-overprovisioned instance family • xn4, n4, mn4, and e4, previous-generation shared instance families

Arm-based enterprise-level computing instance families

YiTian 710 processor-powered instance families	**Ampere® Altra® processor**-powered instance families
• g8y, general-purpose instance family • c8y, compute-optimized instance family • r8y, memory-optimized instance family	• g6r, general-purpose instance family • c6r, compute-optimized instance family

ECS Bare Metal Instance families

Recommended instance families					Not recommended instance families (If the following instance families are sold out, we recommend that you use the instance families in the preceding columns.)
General-purpose instance families (g series)	**Compute-optimized instance families (c series)**	**Memory-optimized instance families (r series)**	**Instance families with high clock speeds (hf series)**	**GPU-accelerated compute-optimized instance families (gn series)**	
• ebmg8i, general-purpose ECS Bare Metal Instance family • ebmg7, general-purpose ECS Bare Metal Instance family • ebmg7a, general-purpose ECS Bare Metal Instance family • ebmg6a, general-purpose ECS Bare Metal Instance family • ebmg6e, performance-enhanced general-purpose ECS Bare Metal Instance family • ebmg6, general-purpose ECS Bare Metal Instance family	• ebmc8i, compute-optimized ECS Bare Metal Instance family • ebmc7, compute-optimized ECS Bare Metal Instance family • ebmc7a, compute-optimized ECS Bare Metal Instance family • ebmc6me, compute-optimized ECS Bare Metal Instance family • ebmc6a, compute-optimized ECS Bare Metal Instance family • ebmc6e, performance-enhanced compute-optimized ECS Bare Metal Instance family • ebmc6, compute-optimized ECS Bare Metal Instance family	• ebmr7, memory-optimized ECS Bare Metal Instance family • ebmr7a, memory-optimized ECS Bare Metal Instance family • ebmr6a, memory-optimized ECS Bare Metal Instance family • ebmr6e, performance-enhanced memory-optimized ECS Bare Metal Instance family • ebmr6, memory-optimized ECS Bare Metal Instance family • ebmre6p, performance-enhanced persistent-memory-optimized ECS Bare Metal Instance family • ebmre6-6t, performance-enhanced memory-optimized ECS Bare Metal Instance family	• ebmhfg7, general-purpose ECS Bare Metal Instance family with high clock speeds • ebmhfc7, compute-optimized ECS Bare Metal Instance family with high clock speeds • ebmhfr7, memory-optimized ECS Bare Metal Instance family with high clock speeds • ebmhfg6, general-purpose ECS Bare Metal Instance family with high clock speeds • ebmhfc6, compute-optimized ECS Bare Metal Instance family with high clock speeds • ebmhfr6, memory-optimized ECS Bare Metal Instance family with high clock speeds	• ebmgn8v, GPU-accelerated compute-optimized ECS Bare Metal Instance family • ebmgn8is, GPU-accelerated compute-optimized ECS Bare Metal Instance family • ebmgn7e, GPU-accelerated compute-optimized ECS Bare Metal Instance family • ebmgn7i, GPU-accelerated compute-optimized ECS Bare Metal Instance family • ebmgn7, GPU-accelerated compute-optimized ECS Bare Metal Instance family • ebmgn6ia, GPU-accelerated compute-optimized ECS Bare Metal Instance family	• ebmc5s, network-enhanced compute-optimized ECS Bare Metal Instance family • ebmg5s, network-enhanced general-purpose ECS Bare Metal Instance family • ebmr5s, network-enhanced memory-optimized ECS Bare Metal Instance family • ebmg5, general-purpose ECS Bare Metal Instance family

				• ebmgn6e, GPU-accelerated compute-optimized ECS Bare Metal Instance family • ebmgn6v, GPU-accelerated compute-optimized ECS Bare Metal Instance family • ebmgn6i, GPU-accelerated compute-optimized ECS Bare Metal Instance family	

Super Computing Cluster (SCC) instance families

- scchfc6, compute-optimized SCC instance family with high clock speeds
- scchfg6, general-purpose SCC instance family with high clock speeds
- scchfr6, memory-optimized SCC instance family with high clock speeds
- scch5, SCC instance family with high clock speeds

Enterprise-level heterogeneous computing instance families

Recommended instance families	Not recommended instance families (If the following instance families are sold out, we recommend that you use the recommended instance families.)
• sgn7i-vws, vGPU-accelerated instance family with shared CPUs • vgn7i-vws, vGPU-accelerated instance family • vgn6i-vws, vGPU-accelerated instance family • gn8v, GPU-accelerated compute-optimized instance family • gn8is, GPU-accelerated compute-optimized instance family • gn7e, GPU-accelerated compute-optimized instance family • gn7i, GPU-accelerated compute-optimized instance family • gn7s, GPU-accelerated compute-optimized instance family • gn7, GPU-accelerated compute-optimized instance family • gn6i, GPU-accelerated compute-optimized instance family • gn6e, GPU-accelerated compute-optimized instance family • gn6v, GPU-accelerated compute-optimized instance family	• gn5, GPU-accelerated compute-optimized instance family • gn5i, GPU-accelerated compute-optimized instance family

x86-based enterprise-level computing instance families

g8a, general-purpose instance family

- **Introduction**: This instance family uses the innovative Cloud Infrastructure Processing Unit (CIPU) architecture developed by Alibaba Cloud to provide consistent computing power, a more robust I/O engine, and chip-level security hardening.
- **Supported scenarios**: general-purpose enterprise-level applications such as Java, in-memory database and relational database applications, big data applications such as Kafka and Elasticsearch, web applications, AI training and inference, and audio and video transcoding applications.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.7 GHz AMD EPYC™ Genoa 9T24 processors that deliver a turbo frequency of up to 3.7 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports the Non-Volatile Memory Express (NVMe) protocol. For more information, see NVMe protocol.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports elastic RDMA interfaces (ERIs). For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides ultra-high packet forwarding rates.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.
- **Security**: Supports the virtual Trusted Platform Module (vTPM) feature. For more information, see Overview.

g8a instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)

ecs.g8a.large	2	8	1.5/12.5	900,000	Up to 250,000	2	3	6	6	20,000/110,000	1.5/10
ecs.g8a.xlarge	4	16	2.5/12.5	1,000,000	Up to 250,000	4	4	6	6	30,000/110,000	2/10
ecs.g8a.2xlarge	8	32	4/12.5	1,600,000	Up to 250,000	8	4	15	15	45,000/110,000	2.5/10
ecs.g8a.4xlarge	16	64	7/12.5	2,000,000	300,000	16	8	30	30	60,000/110,000	3.5/10
ecs.g8a.8xlarge	32	128	10/25	3,000,000	600,000	32	8	30	30	80,000/110,000	5/10
ecs.g8a.12xlarge	48	192	16/25	4,500,000	750,000	48	8	30	30	120,000/none	8/10
ecs.g8a.16xlarge	64	256	20/25	6,000,000	1,000,000	64	8	30	30	160,000/none	10/none
ecs.g8a.24xlarge	96	384	32/none	9,000,000	1,500,000	64	15	30	30	240,000/none	16/none
ecs.g8a.32xlarge	128	512	40/none	12,000,000	2,000,000	64	15	30	30	320,000/none	20/none
ecs.g8a.48xlarge	192	768	64/none	18,000,000	3,000,000	64	15	30	30	500,000/none	32/none

Note

- Packet forwarding rates significantly vary based on business scenarios. We recommend that you perform business stress tests on instances to select appropriate instance types.
- For ecs.g8a.large and ecs.g8a.xlarge instances, you must enable the Jumbo Frames feature before the instances can burst their network bandwidths to 12.5 Gbit/s. For more information, see Jumbo Frames.

g8i, general-purpose instance family

- **Introduction**: This instance family uses the innovative CIPU architecture developed by Alibaba Cloud to provide consistent computing power, a more robust I/O engine, and chip-level security hardening.

- **Supported scenarios**: scenarios where large volumes of packets are received and transmitted, game servers, small and medium-sized database systems, caches, search clusters, search promotion applications, websites, application servers, data analytics and computing, and scenarios that require secure and trusted computing.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses Intel® Xeon® Emerald Rapids or Intel® Xeon® Sapphire Rapids processors that deliver a clock speed of at least 2.7 GHz and an all-core turbo frequency of 3.2 GHz to provide consistent computing performance.

    > **Note**
    >
    > When you purchase an instance of this instance family, the system randomly allocates one type of the preceding processors to the instance. You cannot select a processor type for the instance.

  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports the Non-Volatile Memory Express (NVMe) protocol. For more information, see NVMe protocol.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports elastic RDMA interfaces (ERIs). For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides high network performance based on large computing capacity.
- **Security**:
  - Supports the vTPM feature. For more information, see Overview.
  - Supports Intel Total Memory Encryption (TME) to encrypt memory.

g8i instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.g8i.large	2	8	2.5/burstable up to 15	1,000,000	Up to 300,000	2	3	6	6	25,000/burstable up to 200,000	2/burstable up to 10

ecs.g8i.xlarge	4	16	4/burstable up to 15	1,200,000	Up to 300,000	4	4	15	15	50,000/burstable up to 200,000	2.5/burstable up to 10
ecs.g8i.2xlarge	8	32	6/burstable up to 15	1,600,000	Up to 300,000	8	4	15	15	60,000/burstable up to 200,000	4/burstable up to 10
ecs.g8i.3xlarge	12	48	10/burstable up to 15	2,400,000	Up to 300,000	12	8	15	15	80,000/burstable up to 200,000	5/burstable up to 10
ecs.g8i.4xlarge	16	64	12/burstable up to 25	3,000,000	350,000	16	8	30	30	100,000/burstable up to 200,000	6/burstable up to 10
ecs.g8i.6xlarge	24	96	15/burstable up to 25	4,500,000	500,000	24	8	30	30	120,000/burstable up to 200,000	7.5/burstable up to 10
ecs.g8i.8xlarge	32	128	20/burstable up to 25	6,000,000	800,000	32	8	30	30	200,000/none	10/none
ecs.g8i.12xlarge	48	192	25/none	9,000,000	1,000,000	48	8	30	30	300,000/none	12/none
ecs.g8i.16xlarge	64	256	32/none	12,000,000	1,600,000	64	8	30	30	360,000/none	20/none
ecs.g8i.24xlarge	96	384	50/none	18,000,000	2,000,000	64	15	30	30	500,000/none	24/none
ecs.g8i.48xlarge	192	1,024	100/none	30,000,000	4,000,000	64	15	50	50	1,000,000/none	48/none

> **Note**
>
> If you want to use the ecs.g8i.48xlarge instance type, submit a ticket.

g8ae, performance-enhanced general-purpose instance family

- **Introduction**: This instance family uses the innovative CIPU architecture developed by Alibaba Cloud to provide consistent computing power, a more robust I/O engine, and chip-level security hardening.
- **Supported scenarios**: AI scenarios such as deep learning, training, and AI inference, high-performance scientific computing scenarios such as high-performance computing (HPC), large and medium-sized database systems, caches, search clusters, servers for massively multiplayer online (MMO) games, and other general-purpose enterprise-level applications that require high performance.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 3.4 GHz AMD EPYC™ Genoa processors that deliver a single-core turbo frequency of up to 3.75 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports the Non-Volatile Memory Express (NVMe) protocol. For more information, see NVMe protocol.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports elastic RDMA interfaces (ERIs). For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides ultra-high packet forwarding rates.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.
- **Security**: Supports the virtual Trusted Platform Module (vTPM) feature. For more information, see Overview.

g8ae instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Support for vTPM	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.g8ae.large	2	8	3/burstable up to 15	1,000,000	Yes	Up to 300,000	2	3	6	6	30,000/burstable up to 200,000	2/burstable up to 10
ecs.g8ae.xlarge	4	16	4/burstable up to 15	1,200,000	Yes	Up to 300,000	4	4	15	15	50,000/burstable up to 200,000	2.5/burstable up to 10
ecs.g8ae.2xlarge	8	32	6/burstable up to 15	1,600,000	Yes	Up to 300,000	8	4	15	15	60,000/burstable up to 200,000	3/burstable up to 10
ecs.g8ae.4xlarge	16	64	12/burstable up to 25	3,000,000	Yes	500,000	16	8	30	30	100,000/burstable up to 200,000	6/burstable up to 10
ecs.g8ae.8xlarge	32	128	20/burstable up to 25	6,000,000	Yes	1,000,000	32	8	30	30	200,000/none	10/none
ecs.g8ae.16xlarge	64	256	32/none	9,000,000	Yes	1,500,000	64	8	30	30	250,000/none	16/none
ecs.g8ae.32xlarge	128	512	64/none	18,000,000	Yes	3,000,000	64	15	30	30	500,000/none	32/none

> **Note**
>
> For ecs.g8ae.large and ecs.g8ae.xlarge instances, you must enable the Jumbo Frames feature before the instances can burst their network bandwidths to 15 Gbit/s. For more information, see Jumbo Frames.

g7a, general-purpose instance family

- **Introduction**: This instance family uses the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**: video encoding and decoding, scenarios where large volumes of packets are received and transmitted, websites, application servers, small and medium-sized database systems, caches, search clusters, game servers, scenarios where applications such as DevOps applications are developed and tested, and other general-purpose enterprise-level applications.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.55 GHz AMD EPYC™ MILAN processors that deliver a single-core turbo frequency of up to 3.5 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.

g7a instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.g7a.large	2	8	1/burstable up to 10	900,000	Up to 250,000	2	3	6	6	12,500/burstable up to 110,000	1/burstable up to 6

ecs.g7a.xlarge	4	16	1.5/burstable up to 10	1,000,000	Up to 250,000	4	4	15	15	20,000/burstable up to 110,000	1.5/burstable up to 6
ecs.g7a.2xlarge	8	32	2.5/burstable up to 10	1,600,000	Up to 250,000	8	4	15	15	30,000/burstable up to 110,000	2/burstable up to 6
ecs.g7a.4xlarge	16	64	5/burstable up to 10	2,000,000	300,000	8	8	30	30	60,000/burstable up to 110,000	3.7/burstable up to 10.5
ecs.g7a.8xlarge	32	128	8/burstable up to 10	3,000,000	600,000	16	7	30	30	75,000/burstable up to 110,000	4.1/burstable up to 11
ecs.g7a.16xlarge	64	256	16/none	6,000,000	1,000,000	32	8	30	30	150,000/none	8.2/none
ecs.g7a-nps1.16xlarge	64	256	16/none	6,000,000	1,000,000	32	8	30	30	150,000/none	8.2/none
ecs.g7a.32xlarge	128	512	32/none	12,000,000	2,000,000	32	15	30	30	300,000/none	16.4/none

Note

Ubuntu 16 and Debian 9 operating system kernels do not support AMD EPYC™ MILAN processors. Do not use Ubuntu 16 or Debian 9 images to create instances of this instance family. Instances of this instance family created from Ubuntu 16 or Debian 9 images cannot be started.

g7, general-purpose instance family

- **Introduction**: This instance family uses the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**: scenarios where large volumes of packets are received and transmitted such as live commenting on videos and telecom data forwarding, game servers, small and medium-sized database systems, caches, search clusters, enterprise-level applications of various types and sizes, websites, application servers, data analytics and computing, scenarios that require secure and trusted computing, and blockchain scenarios.

- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses third-generation Intel® Xeon® Scalable (Ice Lake) processors that deliver a base frequency of 2.7 GHz and an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides ultra-high packet forwarding rates.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.
- **Security**:
  - Supports the vTPM feature. For more information, see Overview.
  - Supports the Enclave feature and provides virtualization-based confidential computing environments. For more information, see Build a confidential computing environment by using Enclave.

g7 instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Support for vTPM	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Maximum attached data disks	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.g7.large	2	8	2/burstable up to 12.5	1,100,000	Yes	Up to 500,000	2	3	6	6	8	20,000/burstable up to 160,000	1.5/burstable up to 10
ecs.g7.xlarge	4	16	3/burstable up to 12.5	1,100,000	Yes	Up to 500,000	4	4	15	15	8	40,000/burstable up to 160,000	2/burstable up to 10

ecs.g7.2xlarge	8	32	5/burstable up to 15	1,600,000	Yes	Up to 500,000	8	4	15	15	16	50,000/burstable up to 160,000	3/burstable up to 10
ecs.g7.3xlarge	12	48	8/burstable up to 15	2,400,000	Yes	Up to 500,000	8	8	15	15	16	70,000/burstable up to 160,000	4/burstable up to 10
ecs.g7.4xlarge	16	64	10/burstable up to 25	3,000,000	Yes	500,000	8	8	30	30	16	80,000/burstable up to 160,000	5/burstable up to 10
ecs.g7.6xlarge	24	96	12/burstable up to 25	4,500,000	Yes	550,000	12	8	30	30	16	110,000/burstable up to 160,000	6/10
ecs.g7.8xlarge	32	128	16/burstable up to 32	6,000,000	Yes	600,000	16	8	30	30	24	160,000/none	10/none
ecs.g7.16xlarge	64	256	32/none	12,000,000	Yes	1,200,000	32	8	30	30	32	360,000/none	16/none
ecs.g7.32xlarge	128	512	64/none	24,000,000	Yes	2,400,000	32	15	30	30	32	600,000/none	32/none

g6, general-purpose instance family

- **Introduction**: This instance family offloads a large number of virtualization features to dedicated hardware by using the SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads.
- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Enterprise-level applications of various types and sizes

- Websites and application servers
- Game servers
- Small and medium-sized database systems, caches, and search clusters
- Data analytics and computing
- Computing clusters and memory-intensive data processing

- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver a turbo frequency of 3.2 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.

    > **Note**
    >
    > The maximum performance of disks varies based on the instance families. A single instance of this instance family can deliver up to 200,000 IOPS.

  - Provides high network and storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.

    > **Note**
    >
    > Network performance varies based on the instance families. For higher concurrent connection and network packet forwarding capabilities, we recommend that you use the g7ne instance family.

  - Provides high network performance based on large computing capacity.
- **Supported instance type changes**: Supports changes to c6 or r6 instance types.

g6 instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.g6.large	2	8	1/burstable up to 3	300,000	Up to 250,000	2	2	6	1	10,000	1

ecs.g6.xlarge	4	16	1.5/burstable up to 5	500,000	Up to 250,000	4	3	10	1	20,000	1.5
ecs.g6.2xlarge	8	32	2.5/burstable up to 8	800,000	Up to 250,000	8	4	10	1	25,000	2
ecs.g6.3xlarge	12	48	4/burstable up to 10	900,000	Up to 250,000	8	6	10	1	30,000	2.5
ecs.g6.4xlarge	16	64	5/burstable up to 10	1,000,000	300,000	8	8	20	1	40,000	3
ecs.g6.6xlarge	24	96	7.5/burstable up to 10	1,500,000	450,000	12	8	20	1	50,000	4
ecs.g6.8xlarge	32	128	10/none	2,000,000	600,000	16	8	20	1	60,000	5
ecs.g6.13xlarge	52	192	12.5/none	3,000,000	900,000	32	7	20	1	100,000	8
ecs.g6.26xlarge	104	384	25/none	6,000,000	1,800,000	32	15	20	1	200,000	16

g6a, general-purpose instance family

- **Introduction**: This instance family offloads a large number of virtualization features to dedicated hardware by using the SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads.
- **Supported scenarios**:
  - Video encoding and decoding
  - Scenarios where large volumes of packets are received and transmitted
  - Websites and application servers
  - Small and medium-sized database systems, caches, and search clusters
  - Game servers
  - Scenarios where applications such as DevOps applications are developed and tested
  - Other general-purpose enterprise-level applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.6 GHz AMD EPYC™ ROME processors that deliver a turbo frequency of 3.3 GHz to provide consistent computing performance.

- Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- Is compatible with specific operating systems. For more information, see Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
  - Provides high network and storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

g6a instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.g6a.large	2	8	1/10	900,000	Up to 250,000	2	2	6	1	12,500	1
ecs.g6a.xlarge	4	16	1.5/10	1,000,000	Up to 250,000	4	3	15	1	20,000	1.5
ecs.g6a.2xlarge	8	32	2.5/10	1,600,000	Up to 250,000	8	4	15	1	30,000	2
ecs.g6a.4xlarge	16	64	5/10	2,000,000	300,000	8	8	30	1	60,000	3.1
ecs.g6a.8xlarge	32	128	8/10	3,000,000	600,000	16	7	30	1	75,000	4.1
ecs.g6a.16xlarge	64	256	16/none	6,000,000	1,000,000	32	8	30	1	150,000	8.2
ecs.g6a.32xlarge	128	512	32/none	12,000,000	2,000,000	32	15	30	1	300,000	16.4

g6e, performance-enhanced general-purpose instance family

- **Introduction**: This instance family offloads a large number of virtualization features to dedicated hardware by using the SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Enterprise-level applications of various types and sizes
  - Websites and application servers
  - Game servers
  - Small and medium-sized database systems, caches, and search clusters
  - Data analytics and computing
  - Computing clusters and memory-intensive data processing
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver a turbo frequency of 3.2 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Provides high network and storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.

    > **Note**
    >
    > Network performance varies based on the instance families. For higher concurrent connection and network packet forwarding capabilities, we recommend that you use the g7ne instance family.

  - Provides high network performance based on large computing capacity.

g6e instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)

ecs.g6e.large	2	8	1.2/burstable up to 10	900,000	Up to 250,000	2	3	6	1	20,000	1
ecs.g6e.xlarge	4	16	2/burstable up to 10	1,000,000	Up to 250,000	4	4	15	1	40,000	1.5
ecs.g6e.2xlarge	8	32	3/burstable up to 10	1,600,000	Up to 250,000	8	4	15	1	50,000	2
ecs.g6e.4xlarge	16	64	6/burstable up to 10	3,000,000	300,000	8	8	30	1	80,000	3
ecs.g6e.8xlarge	32	128	10/none	6,000,000	600,000	16	8	30	1	150,000	5
ecs.g6e.13xlarge	52	192	16/none	9,000,000	1,000,000	32	7	30	1	240,000	8
ecs.g6e.26xlarge	104	384	32/none	24,000,000	1,800,000	32	15	30	1	480,000	16

Note

- The results for network capabilities are the maximum values obtained from single-item tests. For example, when network bandwidth is tested, no stress tests are performed on the packet forwarding rate or other network metrics.
- If you want to use the ecs.g6e.26xlarge instance type, submit a ticket.

g5, general-purpose instance family

- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Enterprise-level applications of various types and sizes
  - Small and medium-sized database systems, caches, and search clusters
  - Data analytics and computing
  - Computing clusters and memory-intensive data processing
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.

- Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) or 8269CY (Cascade Lake) processors to provide consistent computing performance.

  > **Note**
  >
  > Instances of this instance family may be deployed on different server platforms. If your business requires all instances to be deployed on the same server platform, we recommend that you use the g6, g6e, or g7 instance family instead.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.

    > **Note**
    >
    > The maximum performance of disks varies based on the instance families. A single instance of this instance family can deliver up to 200,000 IOPS.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.

    > **Note**
    >
    > Network performance varies based on the instance families. For higher concurrent connection and network packet forwarding capabilities, we recommend that you use the g7ne instance family.

  - Provides high network performance based on large computing capacity.

g5 instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.g5.large	2	8	1	300,000	2	2	6	1
ecs.g5.xlarge	4	16	1.5	500,000	2	3	10	1
ecs.g5.2xlarge	8	32	2.5	800,000	4	4	10	1
ecs.g5.3xlarge	12	48	4	900,000	4	6	10	1
ecs.g5.4xlarge	16	64	5	1,000,000	4	8	20	1
ecs.g5.6xlarge	24	96	7.5	1,500,000	6	8	20	1

ecs.g5.8xlarge	32	128	10	2,000,000	8	8	20	1
ecs.g5.16xlarge	64	256	20	4,000,000	16	8	20	1

> **Note**
>
> - You can go to the Instance Types Available for Each Region page to view the instance types available in each region.
> - For more information about these specifications, see the "Instance type specifications" section in Overview of instance families. Packet forwarding rates vary significantly based on business scenarios. We recommend that you perform business stress tests on instances to choose appropriate instance types.

sn2ne, network-enhanced general-purpose instance family

- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Enterprise-level applications of various types and sizes
  - Small and medium-sized database systems, caches, and search clusters
  - Data analytics and computing
  - Computing clusters and memory-intensive data processing
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.5 GHz Intel® Xeon® E5-2682 v4 (Broadwell), Platinum 8163 (Skylake), or 8269CY (Cascade Lake) processors to provide consistent computing performance.

    > **Note**
    >
    > Instances of this instance family may be deployed on different server platforms. If your business requires all instances to be deployed on the same server platform, we recommend that you use the g6, g6e, or g7 instance family instead.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

sn2ne instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.sn2ne.large	2	8	1	300,000	2	2	6	1
ecs.sn2ne.xlarge	4	16	1.5	500,000	2	3	10	1
ecs.sn2ne.2xlarge	8	32	2	1,000,000	4	4	10	1
ecs.sn2ne.3xlarge	12	48	2.5	1,300,000	4	6	10	1
ecs.sn2ne.4xlarge	16	64	3	1,600,000	4	8	20	1
ecs.sn2ne.6xlarge	24	96	4.5	2,000,000	6	8	20	1
ecs.sn2ne.8xlarge	32	128	6	2,500,000	8	8	20	1
ecs.sn2ne.14xlarge	56	224	10	4,500,000	14	8	20	1

c8a, compute-optimized instance family

- **Introduction**: This instance family uses the innovative Cloud Infrastructure Processing Unit (CIPU) architecture developed by Alibaba Cloud to provide consistent computing power, a more robust I/O engine, and chip-level security hardening.
- **Supported scenarios**: big data applications, web applications, AI training and inference, and audio and video transcoding applications.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.7 GHz AMD EPYC™ Genoa processors that deliver a turbo frequency of up to 3.7 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports the Non-Volatile Memory Express (NVMe) protocol. For more information, see NVMe protocol.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.

- Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports elastic RDMA interfaces (ERIs). For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides ultra-high packet forwarding rates.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.
- **Security**: Supports the virtual Trusted Platform Module (vTPM) feature. For more information, see Overview.

The c8a instance family includes the following instance types: ecs.c8a.large, ecs.c8a.xlarge, ecs.c8a.2xlarge, ecs.c8a.4xlarge, ecs.c8a.8xlarge, ecs.c8a.12xlarge, ecs.c8a.16xlarge, ecs.c8a.24xlarge, ecs.c8a.32xlarge, and ecs.c8a.48xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.c8a.large	2	4	1.5/burstable up to 12.5	900,000	Up to 250,000	2	3	6	6	20,000/burstable up to 110,000	1.5/burstable up to 10
ecs.c8a.xlarge	4	8	2.5/burstable up to 12.5	1,000,000	Up to 250,000	4	4	6	6	30,000/burstable up to 110,000	2/burstable up to 10
ecs.c8a.2xlarge	8	16	4/burstable up to 12.5	1,600,000	Up to 250,000	8	4	15	15	45,000/burstable up to 110,000	2.5/burstable up to 10
ecs.c8a.4xlarge	16	32	7/burstable up to 12.5	2,000,000	300,000	16	8	30	30	60,000/burstable up to 110,000	3.5/burstable up to 10

ecs.c8a.8xlarge	32	64	10/burstable up to 25	3,000,000	600,000	32	8	30	30	80,000/burstable up to 110,000	5/burstable up to 10
ecs.c8a.12xlarge	48	96	16/25	4,500,000	750,000	48	8	30	30	120,000/none	8/burstable up to 10
ecs.c8a.16xlarge	64	128	20/25	6,000,000	1,000,000	64	8	30	30	160,000/none	10/none
ecs.c8a.24xlarge	96	192	32/none	9,000,000	1,500,000	64	15	30	30	240,000/none	16/none
ecs.c8a.32xlarge	128	256	40/none	12,000,000	2,000,000	64	15	30	30	320,000/none	20/none
ecs.c8a.48xlarge	192	384	64/none	18,000,000	3,000,000	64	15	30	30	500,000/none	32/none

> **Note**
>
> For ecs.c8a.large and ecs.c8a.xlarge instances, you must enable the Jumbo Frames feature before the instances can burst their network bandwidths to 12.5 Gbit/s. For more information, see Jumbo Frames.

c8i, compute-optimized instance family

- **Introduction**: This instance family uses the innovative CIPU architecture developed by Alibaba Cloud to provide consistent computing power, a more robust I/O engine, and chip-level security hardening.
- **Supported scenarios**: machine learning inference applications, data analytics, batch computing, video encoding, frontend servers for games, high-performance scientific and engineering applications, and web frontend servers.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses Intel® Xeon® Emerald Rapids or Intel® Xeon® Sapphire Rapids processors that deliver a clock speed of at least 2.7 GHz and an all-core turbo frequency of 3.2 GHz to provide consistent computing performance.

    > **Note**
    >
    > When you purchase an instance of this instance family, the system randomly allocates one type of the preceding processors to the instance. You cannot select a processor type for the instance.

  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports the NVMe protocol. For more information, see NVMe protocol.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports elastic RDMA interfaces (ERIs). For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.
- **Security**:
  - Supports the virtual Trusted Platform Module (vTPM) feature. For more information, see Overview.
  - Implements trusted boot based on Trusted Cryptography Module (TCM) or TPM chips to provide ultra-high security capabilities. During a trusted boot, all modules in the boot chain from the underlying server to the ECS instance are measured and verified.
  - Supports Intel Total Memory Encryption (TME) to encrypt memory.

The c8i instance family includes the following instance types: ecs.c8i.large, ecs.c8i.xlarge, ecs.c8i.2xlarge, ecs.c8i.3xlarge, ecs.c8i.4xlarge, ecs.c8i.6xlarge, ecs.c8i.8xlarge, ecs.c8i.12xlarge, ecs.c8i.16xlarge, ecs.c8i.24xlarge, ecs.c8i.48xlarge, and ecs.ebmc8i.48xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.c8i.large	2	4	2.5/burstable up to 15	1,000,000	Up to 300,000	2	3	6	6	25,000/burstable up to 200,000	2/burstable up to 10
ecs.c8i.xlarge	4	8	4/burstable up to 15	1,200,000	Up to 300,000	4	4	15	15	50,000/burstable up to 200,000	2.5/burstable up to 10
ecs.c8i.2xlarge	8	16	6/burstable up to 15	1,600,000	Up to 300,000	8	4	15	15	60,000/burstable up to 200,000	4/burstable up to 10

ecs.c8i.3xlarge	12	24	10/burstable up to 15	2,400,000	Up to 300,000	12	8	15	15	80,000/burstable up to 200,000	5/burstable up to 10
ecs.c8i.4xlarge	16	32	12/burstable up to 25	3,000,000	350,000	16	8	30	30	100,000/burstable up to 200,000	6/burstable up to 10
ecs.c8i.6xlarge	24	48	15/burstable up to 25	4,500,000	500,000	24	8	30	30	120,000/burstable up to 200,000	7.5/burstable up to 10
ecs.c8i.8xlarge	32	64	20/burstable up to 25	6,000,000	800,000	32	8	30	30	200,000/none	10/none
ecs.c8i.12xlarge	48	96	25/none	9,000,000	1,000,000	48	8	30	30	300,000/none	12/none
ecs.c8i.16xlarge	64	128	32/none	12,000,000	1,600,000	64	8	30	30	360,000/none	20/none
ecs.c8i.24xlarge	96	192	50/none	18,000,000	2,000,000	64	15	30	30	500,000/none	24/none
ecs.c8i.48xlarge	192	512	100/none	30,000,000	4,000,000	64	15	50	50	1,000,000/none	48/none

> **Note**
>
> If you want to use the ecs.c8i.48xlarge instance type, submit a ticket.

c8ae, performance-enhanced compute-optimized instance family

- **Introduction**: This instance family uses the innovative CIPU architecture developed by Alibaba Cloud to provide consistent computing power, a more robust I/O engine, and chip-level security hardening.
- **Supported scenarios**:
  - AI scenarios, such as deep learning and training, and AI inference
  - High-performance scientific computing scenarios, such as high-performance computing (HPC)

- Large and medium-sized database systems, caches, and search clusters
- Servers for massively multiplayer online (MMO) games
- Other general-purpose enterprise-level applications that have high performance requirements

- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 3.4 GHz AMD EPYC™ Genoa processors that deliver a single-core turbo frequency of up to 3.75 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports the NVMe protocol. For more information, see NVMe protocol.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports elastic RDMA interfaces (ERIs). For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides ultra-high packet forwarding rates.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.
- **Security**: Supports the virtual Trusted Platform Module (vTPM) feature. For more information, see Overview.

The c8ae instance family includes the following instance types: ecs.c8ae.large, ecs.c8ae.xlarge, ecs.c8ae.2xlarge, ecs.c8ae.4xlarge, ecs.c8ae.8xlarge, ecs.c8ae.16xlarge, and ecs.c8ae.32xlarge. Click the following panel to see a table describing the metrics of instance types. For more information, see Metrics of instance types.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Support for vTPM	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.c8ae.large	2	4	3/burstable up to 15	1,000,000	Yes	Up to 300,000	2	3	6	6	30,000/burstable up to 200,000	2/burstable up to 10

ecs.c8ae.xlarge	4	8	4/burstable up to 15	1,200,000	Yes	Up to 300,000	4	4	15	15	50,000/burstable up to 200,000	2.5/bursta ble up to 10
ecs.c8ae.2xlarge	8	16	6/burstable up to 15	1,600,000	Yes	Up to 300,000	8	4	15	15	60,000/burstable up to 200,000	3/burstable up to 10
ecs.c8ae.4xlarge	16	32	12/burstable up to 25	3,000,000	Yes	500,000	16	8	30	30	100,000/burstable up to 200,000	6/burstable up to 10
ecs.c8ae.8xlarge	32	64	20/burstable up to 25	6,000,000	Yes	1,000,000	32	8	30	30	200,000/none	10/none
ecs.c8ae.16xlarge	64	128	32/none	9,000,000	Yes	1,500,000	64	8	30	30	250,000/none	16/none
ecs.c8ae.32xlarge	128	256	64/none	18,000,000	Yes	3,000,000	64	15	30	30	500,000/none	32/none

Note

For ecs.c8ae.large and ecs.c8ae.xlarge instances, you must enable the Jumbo Frames feature before the instances can burst their network bandwidths to 15 Gbit/s. For more information, see Jumbo Frames.

c7a, compute-optimized instance family

- **Introduction**: This instance family uses the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**:
  - Video encoding and decoding
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Web frontend servers
  - Frontend servers for MMO games

- Scenarios where applications such as DevOps applications are developed and tested
- Data analytics and batch computing
- High-performance scientific and engineering applications
- Enterprise-level applications of various types and sizes

- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.55 GHz AMD EPYC™ MILAN processors that deliver a single-core turbo frequency of up to 3.5 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.

The c7a instance family includes the following instance types: ecs.c7a.large, ecs.c7a.xlarge, ecs.c7a.2xlarge, ecs.c7a.4xlarge, ecs.c7a.8xlarge, ecs.c7a-nps1.8xlarge, ecs.c7a.16xlarge, ecs.c7a-nps1.16xlarge, and ecs.c7a.32xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.c7a.large	2	4	1/burstable up to 10	900,000	Up to 250,000	2	3	6	6	12,500/burstable up to 110,000	1/burstable up to 6
ecs.c7a.xlarge	4	8	1.5/burstable up to 10	1,000,000	Up to 250,000	4	4	15	15	20,000/burstable up to 110,000	1.5/burstable up to 6

ecs.c7a.2xlarge	8	16	2.5/burstable up to 10	1,600,000	Up to 250,000	8	4	15	15	30,000/burstable up to 110,000	2/burstable up to 6
ecs.c7a.4xlarge	16	32	5/burstable up to 10	2,000,000	300,000	8	8	30	30	60,000/burstable up to 110,000	3/burstable up to 6
ecs.c7a.8xlarge	32	64	8/burstable up to 10	3,000,000	600,000	16	7	30	30	75,000/burstable up to 110,000	4/burstable up to 6
ecs.c7a-nps1.8xlarge	32	64	8/burstable up to 10	3,000,000	600,000	16	7	30	30	75,000/burstable up to 110,000	4/burstable up to 6
ecs.c7a.16xlarge	64	128	16/none	6,000,000	1,000,000	32	7	30	30	150,000/none	8/none
ecs.c7a-nps1.16xlarge	64	128	16/none	6,000,000	1,000,000	32	7	30	30	150,000/none	8/none
ecs.c7a.32xlarge	128	256	32/none	12,000,000	2,000,000	32	15	30	30	300,000/none	16/none

Note

Ubuntu 16 and Debian 9 operating system kernels do not support AMD EPYC™ MILAN processors. Do not use Ubuntu 16 or Debian 9 images to create instances of this instance family. Instances of this instance family created from Ubuntu 16 or Debian 9 images cannot be started.

c7, compute-optimized instance family

- **Introduction**: This instance family uses the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**:

- Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
- Frontend servers for MMO games
- Web frontend servers
- Data analytics, batch computing, and video encoding
- High-performance scientific and engineering applications
- Scenarios that require secure and trusted computing
- Enterprise-level applications of various types and sizes
- Blockchain scenarios

- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses the third-generation Intel® Xeon® Scalable (Ice Lake) processors that deliver a base frequency of 2.7 GHz and an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides ultra-high packet forwarding rates.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.
- **Security**:
  - Supports the vTPM feature. For more information, see Overview.
  - Supports the Enclave feature and provides virtualization-based confidential computing environments. For more information, see Build a confidential computing environment by using Enclave.

The c7 instance family includes the following instance types: ecs.c7.large, ecs.c7.xlarge, ecs.c7.2xlarge, ecs.c7.3xlarge, ecs.c7.4xlarge, ecs.c7.6xlarge, ecs.c7.8xlarge, ecs.c7.16xlarge, and ecs.c7.32xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Support for vTPM	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Maximum attached data disks	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)

ecs.c7.large	2	4	2/burstable up to 12.5	1,100,000	Yes	Up to 500,000	2	3	6	6	8	20,000/burstable up to 160,000	1.5/burstable up to 10
ecs.c7.xlarge	4	8	3/burstable up to 12.5	1,100,000	Yes	Up to 500,000	4	4	15	15	8	40,000/burstable up to 160,000	2/burstable up to 10
ecs.c7.2xlarge	8	16	5/burstable up to 15	1,600,000	Yes	Up to 500,000	8	4	15	15	16	50,000/burstable up to 160,000	3/burstable up to 10
ecs.c7.3xlarge	12	24	8/burstable up to 15	2,400,000	Yes	Up to 500,000	8	8	15	15	16	70,000/burstable up to 160,000	4/burstable up to 10
ecs.c7.4xlarge	16	32	10/burstable up to 25	3,000,000	Yes	500,000	8	8	30	30	16	80,000/burstable up to 160,000	5/burstable up to 10
ecs.c7.6xlarge	24	48	12/burstable up to 25	4,500,000	Yes	550,000	12	8	30	30	16	110,000/burstable up to 160,000	6/10
ecs.c7.8xlarge	32	64	16/burstable up to 32	6,000,000	Yes	600,000	16	8	30	30	24	160,000/none	10/none
ecs.c7.16xlarge	64	128	32/none	12,000,000	Yes	1,200,000	32	8	30	30	32	360,000/none	16/none

ecs.c7.32xlarge	128	256	64/none	24,000,000	Yes	2,400,000	32	15	30	30	32	600,000/none	32/none

c6, compute-optimized instance family

- **Introduction**: This instance family offloads a large number of virtualization features to dedicated hardware by using the SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads.
- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Web frontend servers
  - Frontend servers for MMO games
  - Data analytics, batch computing, and video encoding
  - High-performance scientific and engineering applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver a turbo frequency of 3.2 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.

    > **Note**
    >
    > The maximum performance of disks varies based on the instance families. A single instance of this instance family can deliver up to 200,000 IOPS.

  - Provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.
- **Supported instance type changes**: Supports changes to g6 or r6 instance types.

The c6 instance family includes the following instance types: ecs.c6.large, ecs.c6.xlarge, ecs.c6.2xlarge, ecs.c6.3xlarge, ecs.c6.4xlarge, ecs.c6.6xlarge, ecs.c6.8xlarge, ecs.c6.13xlarge, and ecs.c6.26xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.c6.large	2	4	1/burstable up to 3	300,000	Up to 250,000	2	2	6	1	10,000	1
ecs.c6.xlarge	4	8	1.5/burstable up to 5	500,000	Up to 250,000	4	3	10	1	20,000	1.5
ecs.c6.2xlarge	8	16	2.5/burstable up to 8	800,000	Up to 250,000	8	4	10	1	25,000	2
ecs.c6.3xlarge	12	24	4/burstable up to 10	900,000	Up to 250,000	8	6	10	10	30,000	2.5
ecs.c6.4xlarge	16	32	5/burstable up to 10	1,000,000	300,000	8	8	20	1	40,000	3
ecs.c6.6xlarge	24	48	7.5/burstable up to 10	1,500,000	450,000	12	8	20	1	50,000	4
ecs.c6.8xlarge	32	64	10/none	2,000,000	600,000	16	8	20	1	60,000	5
ecs.c6.13xlarge	52	96	12.5/none	3,000,000	900,000	32	7	20	1	100,000	8
ecs.c6.26xlarge	104	192	25/none	6,000,000	1,800,000	32	15	20	1	200,000	16

c6a, compute-optimized instance family

- **Introduction**: This instance family offloads a large number of virtualization features to dedicated hardware by using the SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads.
- **Supported scenarios**: video encoding and decoding, scenarios in which large volumes of packets are received and transmitted, web frontend servers, frontend servers for MMO games, and scenarios where applications such as DevOps applications are developed and tested.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.6 GHz AMD EPYC™ ROME processors that deliver a turbo frequency of 3.3 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
  - Provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The c6a instance family includes the following instance types: ecs.c6a.large, ecs.c6a.xlarge, ecs.c6a.2xlarge, ecs.c6a.4xlarge, ecs.c6a.8xlarge, ecs.c6a.16xlarge, and ecs.c6a.32xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.c6a.large	2	4	1/10	900,000	Up to 250,000	2	2	6	1	12,500	1
ecs.c6a.xlarge	4	8	1.5/10	1,000,000	Up to 250,000	4	3	15	1	20,000	1.5
ecs.c6a.2xlarge	8	16	2.5/10	1,600,000	Up to 250,000	8	4	15	1	30,000	2

ecs.c6a.4xlarge	16	32	5/10	2,000,000	300,000	8	8	30	1	60,000	3.1
ecs.c6a.8xlarge	32	64	8/10	3,000,000	600,000	16	7	30	1	75,000	4.1
ecs.c6a.16xlarge	64	128	16/none	6,000,000	1,000,000	32	8	30	1	150,000	8.2
ecs.c6a.32xlarge	128	256	32/none	12,000,000	2,000,000	32	15	30	1	300,000	16.4

c6e, performance-enhanced compute-optimized instance family

- **Introduction**: This instance family offloads a large number of virtualization features to dedicated hardware by using the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Web frontend servers
  - Frontend servers for MMO games
  - Data analytics, batch computing, and video encoding
  - High-performance scientific and engineering applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver a turbo frequency of 3.2 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.

> **Note**
>
> Network performance varies based on the instance families. For higher concurrent connection and network packet forwarding capabilities, we recommend that you use the g7ne instance family.

- Provides high network performance based on large computing capacity.

The c6e instance family includes the following instance types: ecs.c6e.large, ecs.c6e.xlarge, ecs.c6e.2xlarge, ecs.c6e.4xlarge, ecs.c6e.8xlarge, ecs.c6e.13xlarge, and ecs.c6e.26xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.c6e.large	2	4	1.2/burstable up to 10	900,000	Up to 250,000	2	3	6	1	20,000	1
ecs.c6e.xlarge	4	8	2/burstable up to 10	1,000,000	Up to 250,000	4	4	15	1	40,000	1.5
ecs.c6e.2xlarge	8	16	3/burstable up to 10	1,600,000	Up to 250,000	8	4	15	1	50,000	2
ecs.c6e.4xlarge	16	32	6/burstable up to 10	3,000,000	300,000	8	8	30	1	80,000	3
ecs.c6e.8xlarge	32	64	10/none	6,000,000	600,000	16	8	30	1	150,000	5
ecs.c6e.13xlarge	52	96	16/none	9,000,000	1,000,000	32	7	30	1	240,000	8
ecs.c6e.26xlarge	104	192	32/none	24,000,000	1,800,000	32	15	30	1	480,000	16

c5, compute-optimized instance family

- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Web frontend servers
  - Frontend servers for MMO games

- Data analytics, batch computing, and video encoding
- High-performance scientific and engineering applications

- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) or 8269CY (Cascade Lake) processors to provide consistent computing performance.

    > **Note**
    >
    > Instances of this instance family may be deployed on different server platforms. If your business requires all instances to be deployed on the same server platform, we recommend that you use the c6, c6e, or c7 instance family.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.

    > **Note**
    >
    > The maximum performance of disks varies based on the instance families. A single instance of this instance family can deliver up to 200,000 IOPS.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The c5 instance family includes the following instance types: ecs.c5.large, ecs.c5.xlarge, ecs.c5.2xlarge, ecs.c5.3xlarge, ecs.c5.4xlarge, ecs.c5.6xlarge, ecs.c5.8xlarge, and ecs.c5.16xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.c5.large	2	4	1	300,000	2	2	6	1
ecs.c5.xlarge	4	8	1.5	500,000	2	3	10	1
ecs.c5.2xlarge	8	16	2.5	800,000	4	4	10	1
ecs.c5.3xlarge	12	24	4	900,000	4	6	10	1
ecs.c5.4xlarge	16	32	5	1,000,000	4	8	20	1

ecs.c5.6xlarge	24	48	7.5	1,500,000	6	8	20	1
ecs.c5.8xlarge	32	64	10	2,000,000	8	8	20	1
ecs.c5.16xlarge	64	128	20	4,000,000	16	8	20	1

ic5, compute-intensive instance family

- **Supported scenarios**:
  - Web frontend servers
  - Data analytics, batch computing, and video encoding
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Frontend servers for MMO games
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:1.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) or 8269CY (Cascade Lake) processors that deliver an all-core turbo frequency of 2.7 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports only IPv4.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The ic5 instance family includes the following instance types: ecs.ic5.large, ecs.ic5.xlarge, ecs.ic5.2xlarge, ecs.ic5.3xlarge, ecs.ic5.4xlarge, ecs.ic5.6xlarge, ecs.ic5.8xlarge, and ecs.ic5.16xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI
ecs.ic5.large	2	2	1	300,000	2	2	6
ecs.ic5.xlarge	4	4	1.5	500,000	2	3	10
ecs.ic5.2xlarge	8	8	2.5	800,000	2	4	10

ecs.ic5.3xlarge	12	12	4	900,000	4	6	10
ecs.ic5.4xlarge	16	16	5	1,000,000	4	8	20
ecs.ic5.6xlarge	24	24				8	20
ecs.ic5.8xlarge	32	32				8	20
ecs.ic5.16xlarge	64	64				8	20

sn1ne, network-enhanced compute-optimized instance family

- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Web frontend servers
  - Frontend servers for MMO games
  - Data analytics, batch computing, and video encoding
  - High-performance scientific and engineering applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.5 GHz Intel® Xeon® E5-2682 v4 (Broadwell), Platinum 8163 (Skylake), or 8269CY (Cascade Lake) processors to provide consistent computing performance.

    > **Note**
    >
    > Instances of this instance family may be deployed on different server platforms. If your business requires all instances to be deployed on the same server platform, we recommend that you use the c6, c6e, or c7 instance family.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The sn1ne instance family includes the following instance types: ecs.sn1ne.large, ecs.sn1ne.xlarge, ecs.sn1ne.2xlarge, ecs.sn1ne.3xlarge, ecs.sn1ne.4xlarge, ecs.sn1ne.6xlarge, and ecs.sn1ne.8xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.sn1ne.large	2	4	1	300,000	2	2	6	1
ecs.sn1ne.xlarge	4	8	1.5	500,000	2	3	10	1
ecs.sn1ne.2xlarge	8	16	2	1,000,000	4	4	10	1
ecs.sn1ne.3xlarge	12	24	2.5	1,300,000	4	6	10	1
ecs.sn1ne.4xlarge	16	32	3	1,600,000	4	8	20	1
ecs.sn1ne.6xlarge	24	48	4.5	2,000,000	6	8	20	1
ecs.sn1ne.8xlarge	32	64	6	2,500,000	8	8	20	1

r8a, memory-optimized instance family

- **Introduction**: This instance family uses the innovative Cloud Infrastructure Processing Unit (CIPU) architecture developed by Alibaba Cloud to provide stable computing power, a more robust I/O engine, and chip-level security hardening.
- **Supported scenarios**:
  - Memory-intensive, general-purpose, enterprise-level applications such as Java
  - Various in-memory database applications such as Redis and Memcache
  - Big data applications such as Kafka and Elasticsearch
  - Audio and video transcoding applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses 2.7 GHz AMD EPYC™ Genoa processors that deliver a turbo frequency of up to 3.7 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see the Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types section of the "Compatibility between AMD instance types and operating systems" topic.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports the Non-Volatile Memory Express (NVMe) protocol. For more information, see NVMe protocol.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.

- Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports Elastic RDMA interfaces (ERIs). For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides ultra-high packet forwarding rates.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.
- **Security**: Supports the virtual Trusted Platform Module (vTPM) feature. For more information, see Overview of trusted computing capabilities.

The r8a instance family includes the following instance types: ecs.r8a.large, ecs.r8a.xlarge, ecs.r8a.2xlarge, ecs.r8a.4xlarge, ecs.r8a.8xlarge, ecs.r8a.12xlarge, ecs.r8a.16xlarge, ecs.r8a.24xlarge, ecs.r8a.32xlarge, and ecs.r8a.48xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see the Instance type metrics section of the "Classification and naming of instance types" topic.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.r8a.large	2	16	1.5/burstable up to 12.5	900,000	Up to 250,000	2	3	6	6	25,000/burstable up to 110,000	1.5/burstable up to 10
ecs.r8a.xlarge	4	32	2.5/burstable up to 12.5	1,000,000	Up to 250,000	4	4	6	6	30,000/burstable up to 110,000	2/burstable up to 10
ecs.r8a.2xlarge	8	64	4/burstable up to 12.5	1,600,000	Up to 250,000	8	4	15	15	45,000/burstable up to 110,000	2.5/burstable up to 10
ecs.r8a.4xlarge	16	128	7/burstable up to 12.5	2,000,000	300,000	16	8	30	30	60,000/burstable up to 110,000	3.5/burstable up to 10

ecs.r8a.8xlarge	32	256	10/burstable up to 25	3,000,000	600,000	32	8	30	30	80,000/burstable up to 110,000	5/burstable up to 10
ecs.r8a.12xlarge	48	384	16/25	4,500,000	750,000	48	8	30	30	120,000/none	8/burstable up to 10
ecs.r8a.16xlarge	64	512	20/25	6,000,000	1,000,000	64	8	30	30	160,000/none	10/none
ecs.r8a.24xlarge	96	768	32/none	9,000,000	1,500,000	64	15	30	30	240,000/none	16/none
ecs.r8a.32xlarge	128	1024	40/none	12,000,000	2,000,000	64	15	30	30	320,000/none	20/none
ecs.r8a.48xlarge	192	1536	64/none	18,000,000	3,000,000	64	15	30	30	500,000/none	32/none

Note

- For ecs.r8a.large and ecs.r8a.xlarge instances, you must enable the Jumbo Frames feature before the instances can burst their network bandwidths to 12.5 Gbit/s. For more information, see Jumbo Frames.
- To use the ecs.r8a.48xlarge instance type, submit a ticket.

r8i, memory-optimized instance family

- **Introduction**: This instance family uses the innovative CIPU architecture developed by Alibaba Cloud to provide stable computing power, a more robust I/O engine, and chip-level security hardening.
- **Supported scenarios**:
  - Data analytics and mining
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters
  - Distributed in-memory cache, such as Redis
  - Websites and application servers
  - Servers of massively multiplayer online (MMO) games
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses Intel® Xeon® Emerald Rapids or Intel® Xeon® Sapphire Rapids processors that deliver a clock speed of at least 2.7 GHz and an all-core turbo frequency of 3.2 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports the NVMe protocol. For more information, see NVMe protocol.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports ERIs. For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.
- **Security**:
  - Supports the vTPM feature. For more information, see Overview of trusted computing capabilities.
  - Supports Intel Total Memory Encryption (TME) to encrypt memory.

The r8i instance family includes the following instance types: ecs.r8i.large, ecs.r8i.xlarge, ecs.r8i.2xlarge, ecs.r8i.3xlarge, ecs.r8i.4xlarge, ecs.r8i.6xlarge, ecs.r8i.8xlarge, ecs.r8i.12xlarge, ecs.r8i.16xlarge, and ecs.r8i.32xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see the Instance type metrics section of the "Classification and naming of instance types" topic.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.r8i.large	2	16	2.5/burstable up to 15	1,000,000	Up to 300,000	2	3	6	6	25,000/burstable up to 200,000	2/burstable up to 10
ecs.r8i.xlarge	4	32	4/burstable up to 15	1,200,000	Up to 300,000	4	4	15	15	50,000/burstable up to 200,000	2.5/burstable up to 10
ecs.r8i.2xlarge	8	64	6/burstable up to 15	1,600,000	Up to 300,000	8	4	15	15	60,000/burstable up to 200,000	4/burstable up to 10

ecs.r8i.3xlarge	12	96	10/burstable up to 15	2,400,000	Up to 300,000	12	8	15	15	80,000/burstable up to 200,000	5/burstable up to 10
ecs.r8i.4xlarge	16	128	12/burstable up to 25	3,000,000	350,000	16	8	30	30	100,000/burstable up to 200,000	6/burstable up to 10
ecs.r8i.6xlarge	24	192	15/burstable up to 25	4,500,000	500,000	24	8	30	30	120,000/burstable up to 200,000	7.5/burstable up to 10
ecs.r8i.8xlarge	32	256	20/burstable up to 25	6,000,000	800,000	32	8	30	30	200,000/none	10/none
ecs.r8i.12xlarge	48	384	25/none	9,000,000	1,000,000	48	8	30	30	300,000/none	12/none
ecs.r8i.16xlarge	64	512	32/none	12,000,000	1,600,000	64	8	30	30	360,000/none	20/none
ecs.r8i.32xlarge	128	1024	64/none	24,000,000	3,000,000	64	15	30	30	700,000/none	40/none

Note

To use the ecs.r8i.16xlarge and ecs.r8i.32xlarge instance types, submit a ticket.

r8ae, enhanced-performance memory-optimized instance family

- **Introduction**: This instance family uses the innovative CIPU architecture developed by Alibaba Cloud to provide stable computing power, a more robust I/O engine, and chip-level security hardening.
- **Supported scenarios**:
  - AI scenarios, such as deep learning and training, and AI inference
  - High-performance scientific computing scenarios such as high-performance computing (HPC)
  - Large and medium-sized database systems, caches, and search clusters

- Servers of MMO games
- Other general-purpose enterprise-level applications that have high performance requirements

- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses 3.4 GHz AMD EPYC™ Genoa processors that deliver a single-core turbo frequency of up to 3.75 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see the Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types section of the "Compatibility between AMD instance types and operating systems" topic.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports the NVMe protocol. For more information, see NVMe protocol.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports ERIs. For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides ultra-high packet forwarding rates.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.
- **Security**: Supports the vTPM feature. For more information, see Overview of trusted computing capabilities.

The r8ae instance family includes the following instance types: ecs.r8ae.large, ecs.r8ae.xlarge, ecs.r8ae.2xlarge, ecs.r8ae.4xlarge, and ecs.r8ae.8xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see the Instance type metrics section of the "Classification and naming of instance types" topic.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Support for vTPM	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.r8ae.large	2	16	3/burstable up to 15	1,000,000	Yes	Up to 300,000	2	3	6	6	30,000/burstable up to 200,000	2/burstable up to 10

ecs.r8ae.xlarge	4	32	4/burstable up to 15	1,200,000	Yes	Up to 300,000	4	4	15	15	50,000/burstable up to 200,000	2.5/burstable up to 10
ecs.r8ae.2xlarge	8	64	6/burstable up to 15	1,600,000	Yes	Up to 300,000	8	4	15	15	60,000/burstable up to 200,000	3/burstable up to 10
ecs.r8ae.4xlarge	16	128	12/burstable up to 25	3,000,000	Yes	500,000	16	8	30	30	100,000/burstable up to 200,000	6/burstable up to 10
ecs.r8ae.8xlarge	32	256	20/burstable up to 25	6,000,000	Yes	1,000,000	32	8	30	30	200,000/none	10/none

Note

For ecs.r8ae.large and ecs.r8ae.xlarge instances, you must enable the Jumbo Frames feature before the instances can burst their network bandwidths to 15 Gbit/s. For more information, see Jumbo Frames.

r7a, memory-optimized instance family

- **Introduction**: This instance family uses the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses 2.55 GHz AMD EPYC™ MILAN processors that deliver a single-core turbo frequency of up to 3.5 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see the Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types section of the "Compatibility between AMD instance types and operating systems" topic.
- **Supported scenarios**:
  - High-performance databases and in-memory databases
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters

- Blockchain applications

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.

The r7a instance family includes the following instance types: ecs.r7a.large, ecs.r7a.xlarge, ecs.r7a.2xlarge, ecs.r7a.4xlarge, ecs.r7a.8xlarge, ecs.r7a-nps1.8xlarge, ecs.r7a.16xlarge, ecs.r7a-nps1.16xlarge, and ecs.r7a.32xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see the Instance type metrics section of the "Classification and naming of instance types" topic.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.r7a.large	2	16	1/burstable up to 10	900,000	Up to 250,000	2	3	6	6	12,500/burstable up to 110,000	1/burstable up to 6
ecs.r7a.xlarge	4	32	1.5/burstable up to 10	1,000,000	Up to 250,000	4	4	15	15	20,000/burstable up to 110,000	1.5/burstable up to 6
ecs.r7a.2xlarge	8	64	2.5/burstable up to 10	1,600,000	Up to 250,000	8	4	15	15	30,000/burstable up to 110,000	2/burstable up to 6
ecs.r7a.4xlarge	16	128	5/burstable up to 10	2,000,000	300,000	8	8	30	30	60,000/burstable up to 110,000	3/burstable up to 6

ecs.r7a.8xlarge	32	256	8/burstable up to 10	3,000,000	600,000	16	7	30	30	75,000/burstable up to 110,000	4/burstable up to 6
ecs.r7a-nps1.8xlarge	32	256	8/burstable up to 10	3,000,000	8/burstable up to 10	16	7	30	30	75,000/burstable up to 110,000	4/burstable up to 6
ecs.r7a.16xlarge	64	512	16/none	6,000,000	1,000,000	32	7	30	30	150,000/none	8/none
ecs.r7a-nps1.16xlarge	64	512	16/none	6,000,000	1,000,000	32	7	30	30	150,000/none	8/none
ecs.r7a.32xlarge	128	1024	32/none	12,000,000	2,000,000	32	15	30	30	300,000/none	16/none

Note

Ubuntu 16 and Debian 9 operating system kernels do not support AMD EPYC™ MILAN processors. Do not use Ubuntu 16 or Debian 9 images to create instances of this instance family. Instances of this instance family created from Ubuntu 16 or Debian 9 images cannot be started.

r7, memory-optimized instance family

- **Introduction**: This instance family uses the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**:
  - High-performance databases and in-memory databases
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters
  - Scenarios that require secure and trusted computing
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses third-generation Intel® Xeon® Scalable (Ice Lake) processors that deliver a base frequency of 2.7 GHz and an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.
  - Allows you to enable or disable Hyper-Threading.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.

- Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides ultra-high packet forwarding rates.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.
- **Security**:
  - Supports the vTPM feature. For more information, see Overview of trusted computing capabilities.
  - Supports the Enclave feature and provides a virtualization-based confidential computing environment. For more information, see Build a confidential computing environment by using Enclave.

The r7 instance family includes the following instance types: ecs.r7.large, ecs.r7.xlarge, ecs.r7.2xlarge, ecs.r7.3xlarge, ecs.r7.4xlarge, ecs.r7.6xlarge, ecs.r7.8xlarge, ecs.r7.16xlarge, and ecs.r7.32xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see the Instance type metrics section of the "Classification and naming of instance types" topic.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Support for vTPM	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Maximum attached data disks	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.r7.large	2	16	2/burstable up to 12.5	1,100,000	Yes	Up to 500,000	2	3	6	6	8	20,000/burstable up to 160,000	1.5/burstable up to 10
ecs.r7.xlarge	4	32	3/burstable up to 12.5	1,100,000	Yes	Up to 500,000	4	4	15	15	8	40,000/burstable up to 160,000	2/burstable up to 10

ecs.r7.2xlarge	8	64	5/burstable up to 15	1,600,000	Yes	Up to 500,000	8	4	15	15	16	50,000/burstable up to 160,000	3/burstable up to 10
ecs.r7.3xlarge	12	96	8/burstable up to 15	2,400,000	Yes	Up to 500,000	8	8	15	15	16	70,000/burstable up to 160,000	4/burstable up to 10
ecs.r7.4xlarge	16	128	10/burstable up to 25	3,000,000	Yes	500,000	8	8	30	30	16	80,000/burstable up to 160,000	5/burstable up to 10
ecs.r7.6xlarge	24	192	12/burstable up to 25	4,500,000	Yes	550,000	12	8	30	30	16	110,000/160,000	6/10
ecs.r7.8xlarge	32	256	16/burstable up to 32	6,000,000	Yes	600,000	16	8	30	30	24	160,000/none	10/none
ecs.r7.16xlarge	64	512	32/none	12,000,000	Yes	1,200,000	32	8	30	30	32	360,000/none	20/none
ecs.r7.32xlarge	128	1024	64/none	24,000,000	Yes	2,400,000	32	15	30	30	32	600,000/none	32/none

r6, memory-optimized instance family

- **Introduction**: This instance family offloads a large number of virtualization features to dedicated hardware by using the SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads.
- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - High-performance databases and in-memory databases
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters

- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver a turbo frequency of 3.2 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.

    > **Note**
    >
    > The maximum performance of disks varies based on the instance families. A single instance of this instance family can deliver up to 200,000 IOPS.

  - Provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.
- **Instance type change**: Supports changes to g6 or c6 instance types.

The r6 instance family includes the following instance types: ecs.r6.large, ecs.r6.xlarge, ecs.r6.2xlarge, ecs.r6.3xlarge, ecs.r6.4xlarge, ecs.r6.6xlarge, ecs.r6.8xlarge, ecs.r6.13xlarge, and ecs.r6.26xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see the Instance type metrics section of the "Classification and naming of instance types" topic.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.r6.large	2	16	1/burstable up to 3	300,000	Up to 250,000	2	2	6	1	10,000	1
ecs.r6.xlarge	4	32	1.5/burstable up to 5	500,000	Up to 250,000	4	3	10	1	20,000	1.5
ecs.r6.2xlarge	8	64	2.5/burstable up to 8	800,000	Up to 250,000	8	4	10	1	25,000	2

ecs.r6.3xlarge	12	96	4/burstable up to 10	900,000	Up to 250,000	8	6	10	1	30,000	2.5
ecs.r6.4xlarge	16	128	5/burstable up to 10	1,000,000	300,000	8	8	20	1	40,000	3
ecs.r6.6xlarge	24	192	7.5/burstable up to 10	1,500,000	450,000	12	8	20	1	50,000	4
ecs.r6.8xlarge	32	256	10/none	2,000,000	600,000	16	8	20	1	60,000	5
ecs.r6.13xlarge	52	384	12.5/none	3,000,000	900,000	32	7	20	1	100,000	8
ecs.r6.26xlarge	104	768	25/none	6,000,000	1,800,000	32	15	20	1	200,000	16

r6a, memory-optimized instance family

- **Introduction**: This instance family offloads a large number of virtualization features to dedicated hardware by using the SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads.
- **Supported scenarios**: scenarios where large volumes of packets are received and transmitted, video encoding and decoding, in-memory databases, enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters, and scenarios where applications such as DevOps applications are developed and tested.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses 2.6 GHz AMD EPYC™ ROME processors that deliver a turbo frequency of 3.3 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see the Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types section of the "Compatibility between AMD instance types and operating systems" topic.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
  - Provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.

- Provides high network performance based on large computing capacity.

The r6a instance family includes the following instance types: ecs.r6a.large, ecs.r6a.xlarge, ecs.r6a.2xlarge, ecs.r6a.4xlarge, ecs.r6a.8xlarge, and ecs.r6a.16xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see the Instance type metrics section of the "Classification and naming of instance types" topic.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.r6a.large	2	16	1/10	900,000	Up to 250,000	2	2	6	1	12,500	1
ecs.r6a.xlarge	4	32	1.5/10	1,000,000	Up to 250,000	4	3	15	1	20,000	1.5
ecs.r6a.2xlarge	8	64	2.5/10	1,600,000	Up to 250,000	8	4	15	1	30,000	2
ecs.r6a.4xlarge	16	128	5/10	2,000,000	300,000	8	8	30	1	60,000	3.1
ecs.r6a.8xlarge	32	256	8/10	3,000,000	600,000	16	7	30	1	75,000	4.1
ecs.r6a.16xlarge	64	512	16/none	6,000,000	1,000,000	32	8	30	1	150,000	8.2

r6e, enhanced-performance memory-optimized instance family

- **Introduction**: This instance family offloads a large number of virtualization features to dedicated hardware by using the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - High-performance databases and in-memory databases
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters

- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver a turbo frequency of 3.2 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.

> **Note**
>
> Network performance varies based on the instance families. For higher concurrent connection and network packet forwarding capabilities, we recommend that you use the g7ne instance family.

The r6e instance family includes the following instance types: ecs.r6e.large, ecs.r6e.xlarge, ecs.r6e.2xlarge, ecs.r6e.4xlarge, ecs.r6e.8xlarge, ecs.r6e.13xlarge, and ecs.r6e.26xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see the Instance type metrics section of the "Classification and naming of instance types" topic.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.r6e.large	2	16	1.2/burstable up to 10	900,000	Up to 250,000	2	3	6	1	20,000	1
ecs.r6e.xlarge	4	32	2/burstable up to 10	1,000,000	Up to 250,000	4	4	15	1	40,000	1.5
ecs.r6e.2xlarge	8	64	3/burstable up to 10	1,600,000	Up to 250,000	8	4	15	1	50,000	2
ecs.r6e.4xlarge	16	128	6/burstable up to 10	3,000,000	300,000	8	8	30	1	80,000	3

ecs.r6e.8xlarge	32	256	10/none	6,000,000	600,000	16	8	30	1	150,000	5
ecs.r6e.13xlarge	52	384	16/none	9,000,000	1,000,000	32	7	30	1	240,000	8
ecs.r6e.26xlarge	104	768	32/none	24,000,000	1,800,000	32	15	30	1	480,000	16

r5, memory-optimized instance family

- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - High-performance databases and in-memory databases
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) or Intel® Xeon® Platinum 8269CY (Cascade Lake) processors to provide consistent computing performance.

    > **Note**
    >
    > Instances of this instance family may be deployed on different server platforms. If your business requires all instances to be deployed on the same server platform, we recommend that you use the r6, r6e, or r7 instance family instead.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.

    > **Note**
    >
    > The maximum performance of disks varies based on the instance families. A single instance of this instance family can deliver up to 200,000 IOPS.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The r5 instance family includes the following instance types: ecs.r5.large, ecs.r5.xlarge, ecs.r5.2xlarge, ecs.r5.3xlarge, ecs.r5.4xlarge, ecs.r5.6xlarge, ecs.r5.8xlarge, and ecs.r5.16xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see the Instance type metrics section of the "Classification and naming of instance types" topic.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.r5.large	2	16	1	300,000	2	2	6	1
ecs.r5.xlarge	4	32	1.5	500,000	2	3	10	1
ecs.r5.2xlarge	8	64	2.5	800,000	4	4	10	1
ecs.r5.3xlarge	12	96	4	900,000	4	6	10	1
ecs.r5.4xlarge	16	128	5	1,000,000	4	8	20	1
ecs.r5.6xlarge	24	192	7.5	1,500,000	6	8	20	1
ecs.r5.8xlarge	32	256	10	2,000,000	8	8	20	1
ecs.r5.16xlarge	64	512	20	4,000,000	16	8	20	1

se1ne, network-enhanced memory-optimized instance family

- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - High-performance databases and in-memory databases
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses 2.5 GHz Intel® Xeon® E5-2682 v4 (Broadwell) or Platinum 8163 (Skylake) or 8269CY (Cascade Lake) processors to provide consistent computing performance.

> **Note**
>
> Instances of this instance family may be deployed on different server platforms. If your business requires all instances to be deployed on the same server platform, we recommend that you use the r6, r6e, or r7 instance family instead.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The se1ne instance family includes the following instance types: ecs.se1ne.large, ecs.se1ne.xlarge, ecs.se1ne.2xlarge, ecs.se1ne.3xlarge, ecs.se1ne.4xlarge, ecs.se1ne.6xlarge, ecs.se1ne.8xlarge, and ecs.se1ne.14xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see the Instance type metrics section of the "Classification and naming of instance types" topic.

Instance types

Instance type	vCPU	Memory (GiB)	Network bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.se1ne.large	2	16	1	300,000	2	2	6	1
ecs.se1ne.xlarge	4	32	1.5	500,000	2	3	10	1
ecs.se1ne.2xlarge	8	64	2	1,000,000	4	4	10	1
ecs.se1ne.3xlarge	12	96	2.5	1,300,000	4	6	10	1
ecs.se1ne.4xlarge	16	128	3	1,600,000	4	8	20	1
ecs.se1ne.6xlarge	24	192	4.5	2,000,000	6	8	20	1
ecs.se1ne.8xlarge	32	256	6	2,500,000	8	8	20	1
ecs.se1ne.14xlarge	56	480	10	4,500,000	14	8	20	1

se1, memory-optimized instance family

- **Supported scenarios**:
  - High-performance databases and in-memory databases
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.

- Uses 2.5 GHz Intel® Xeon® E5-2682 v4 (Broadwell) or Platinum 8163 (Skylake) or 8269CY (Cascade Lake) processors to provide consistent computing performance.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports only IPv4.
  - Provides high network performance based on large computing capacity.

The se1 instance family includes the following instance types: ecs.se1.large, ecs.se1.xlarge, ecs.se1.2xlarge, ecs.se1.4xlarge, ecs.se1.8xlarge, and ecs.se1.14xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see the Instance type metrics section of the "Classification and naming of instance types" topic.

Instance types

Instance type	vCPU	Memory (GiB)	Network bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI
ecs.se1.large	2	16	0.5	100,000	1	2	6
ecs.se1.xlarge	4	32	0.8	200,000	1	3	10
ecs.se1.2xlarge	8	64	1.5	400,000	1	4	10
ecs.se1.4xlarge	16	128	3	500,000	2	8	20
ecs.se1.8xlarge	32	256	6	800,000	3	8	20
ecs.se1.14xlarge	56	480	10	1,200,000	4	8	20

u1, universal instance family

Features:

- Compute:
  - Offers the following CPU-to-memory ratios: 1:2, 1:4, and 1:8.
  - Uses Intel® Xeon® Platinum Scalable processors.

> **Note**
>
> u1 instances are randomly deployed to different server platforms during instance creation and may be migrated across server platforms during the lifecycle of the instances. u1 instances use technological capabilities to promote cross-platform compatibility. However, business performance significantly varies between server platforms. If your business requires performance consistency, we recommend that you use g7, c7, or r7 instances.

- Storage:
  - Is an instance family in which all instances are I/O optimized.
  - Supports enhanced SSDs (ESSDs), ESSD Entry disks, and ESSD AutoPL disks.
- Network:
  - Supports IPv4 and IPv6.
  - Supports only virtual private clouds (VPCs).
  - Provides high network performance based on large computing capacity.
- Supported scenarios:
  - Small and medium-sized enterprise-level applications
  - Websites and application servers
  - Data analytics and computing
  - Small and medium-sized database systems, caches, and search clusters

Instance types

Instance type	vCPUs	Memory size (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.u1-c1m1.large	2	2	1	300,000	Up to 250,000	2	2	6	1	10,000	1
ecs.u1-c1m2.large	2	4	1	300,000	Up to 250,000	2	2	6	1	10,000	1
ecs.u1-c1m4.large	2	8	1	300,000	Up to 250,000	2	2	6	1	10,000	1
ecs.u1-c1m8.large	2	16	1	300,000	Up to 250,000	2	2	6	1	10,000	1
ecs.u1-c1m1.xlarge	4	4	1.5	500,000	Up to 250,000	2	3	10	1	20,000	1.5
ecs.u1-c1m2.xlarge	4	8	1.5	500,000	Up to 250,000	2	3	10	1	20,000	1.5
ecs.u1-c1m4.xlarge	4	16	1.5	500,000	Up to 250,000	2	3	10	1	20,000	1.5

ecs.u1-c1m8.xlarge	4	32	1.5	500,000	Up to 250,000	2	3	10	1	20,000	1.5
ecs.u1-c1m1.2xlarge	8	8	2.5	800,000	Up to 250,000	4	4	10	1	25,000	2
ecs.u1-c1m2.2xlarge	8	16	2.5	800,000	Up to 250,000	4	4	10	1	25,000	2
ecs.u1-c1m4.2xlarge	8	32	2.5	800,000	Up to 250,000	4	4	10	1	25,000	2
ecs.u1-c1m8.2xlarge	8	64	2.5	800,000	Up to 250,000	4	4	10	1	25,000	2
ecs.u1-c1m1.3xlarge	12	12	4	900,000	Up to 250,000	4	6	10	1	30,000	2.5
ecs.u1-c1m2.3xlarge	12	24	4	900,000	Up to 250,000	4	6	10	1	30,000	2.5
ecs.u1-c1m4.3xlarge	12	48	4	900,000	Up to 250,000	4	6	10	1	30,000	2.5
ecs.u1-c1m8.3xlarge	12	96	4	900,000	Up to 250,000	4	6	10	1	30,000	2.5
ecs.u1-c1m1.4xlarge	16	16	5	1,000,000	Up to 300,000	4	8	20	1	40,000	3
ecs.u1-c1m2.4xlarge	16	32	5	1,000,000	Up to 300,000	4	8	20	1	40,000	3
ecs.u1-c1m4.4xlarge	16	64	5	1,000,000	Up to 300,000	4	8	20	1	40,000	3
ecs.u1-c1m8.4xlarge	16	128	5	1,000,000	Up to 300,000	4	8	20	1	40,000	3
ecs.u1-c1m1.8xlarge	32	32	10	2,000,000	Up to 300,000	8	8	20	1	60,000	5

ecs.u1-c1m2.8xlarge	32	64	10	2,000,000	Up to 300,000	8	8	20	1	60,000	5
ecs.u1-c1m4.8xlarge	32	128	10	2,000,000	Up to 300,000	8	8	20	1	60,000	5
ecs.u1-c1m8.8xlarge	32	256	10	2,000,000	Up to 300,000	8	8	20	1	60,000	5

Note

- You can go to the Instance Types Available for Each Region page to view the instance types available in each region.
- For information about the specifications of the instance types, see the Instance type specifications section of the "Overview of instance families" topic.
- Exceptions may occur when you deploy Data Plane Development Kit (DPDK) applications on u1 instances. To resolve the issue, replace Userspace I/O (UIO) drivers with Virtual Function I/O (VFIO) drivers. For more information, see Replace UIO drivers with VFIO drivers.
- For frequently asked questions about universal instances, see the sections that are related to u1 instances in Instance FAQ.

d3s, storage-intensive big data instance family

Features:

- This instance family is equipped with 12-TB, large-capacity, high-throughput local SATA HDDs and can provide a maximum network bandwidth of 64 Gbit/s between instances.
- Supported scenarios:
  - Big data computing and storage business scenarios in which services such as Hadoop MapReduce, HDFS, Hive, and HBase are used
  - Machine learning scenarios such as Spark in-memory computing and MLlib
  - Search and log data processing scenarios in which solutions such as Elasticsearch and Kafka are used
- This instance family supports online replacement and hot swapping of damaged disks to prevent instance shutdown.

  If a local disk fails, you receive a system event. You can handle the system event by initiating the process of repairing the damaged disk. For more information, see O&M scenarios and system events for instances equipped with local disks.

  Important

  After you initiate the process of repairing the damaged disk, data stored on the damaged disk cannot be restored.
- Compute:
  - Uses 2.7 GHz Intel® Xeon® Scalable (Ice Lake) processors that deliver an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.
- Storage:
  - Is an instance family in which all instances are I/O optimized.
  - Supports only ESSDs and ESSD AutoPL disks.

- Network:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

d3s instance types

Instance type	vCPUs	Memory size (GiB)	Local storage (GB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Disk baseline/burst bandwidth (Gbit/s)
ecs.d3s.2xlarge	8	32	4 * 11,918	10/burstable up to 15	2,000,000	3/burstable up to 5
ecs.d3s.4xlarge	16	64	8 * 11,918	25/none	3,000,000	5/none
ecs.d3s.8xlarge	32	128	16 * 11,918	40/none	6,000,000	8/none
ecs.d3s.12xlarge	48	192	24 * 11,918	60/none	9,000,000	12/none
ecs.d3s.16xlarge	64	256	32 * 11,918	80/none	12,000,000	16/none

d3c, compute-intensive big data instance family

Features:

- This instance family is equipped with high-capacity and high-throughput local disks and can provide a maximum bandwidth of 40 Gbit/s between instances.
- Supported scenarios:
  - Big data computing and storage business scenarios in which services such as Hadoop MapReduce, HDFS, Hive, and HBase are used
  - Scenarios in which EMR JindoFS and Object Storage Service (OSS) are used in combination to separately store hot and cold data and decouple storage from computing
  - Machine learning scenarios such as Spark in-memory computing and MLlib
  - Search and log data processing scenarios in which solutions such as Elasticsearch and Kafka are used
- This instance family supports online replacement and hot swapping of damaged disks to prevent instance shutdown.

  If a local disk fails, you receive a system event. You can handle the system event by initiating the process of repairing the damaged disk. For more information, see O&M scenarios and system events for instances equipped with local disks.

  > **Important**
  >
  > After you initiate the process of repairing the damaged disk, data stored on the damaged disk cannot be restored.

- Compute:
  - Uses third-generation 2.9 GHz Intel® Xeon® Scalable (Ice Lake) processors that deliver an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.

- Storage:
  - Is an instance family in which all instances are I/O optimized.
  - Supports only ESSDs and ESSD AutoPL disks.
- Network:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

d3c instance types

Instance type	vCPUs	Memory size (GiB)	Local storage (GB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.d3c.3xlarge	14	56.0	1 * 13,743	8/burstable up to 10	1,600,000	40,000/none	3/none
ecs.d3c.7xlarge	28	112.0	2 * 13,743	16/burstable up to 25	2,500,000	50,000/none	4/none
ecs.d3c.14xlarge	56	224.0	4 * 13,743	40/none	5,000,000	100,000/none	8/none

> **Note**
>
> This instance family supports only Linux images. When you create an instance of this instance family, select a Linux image.

d2c, compute-intensive big data instance family

Features:

- This instance family is equipped with high-capacity and high-throughput local SATA HDDs and can provide a maximum bandwidth of 35 Gbit/s between instances.
- Supported scenarios:
  - Big data computing and storage business scenarios in which services such as Hadoop MapReduce, HDFS, Hive, and HBase are used
  - Scenarios in which EMR JindoFS and OSS are used in combination to separately store hot and cold data and decouple storage from computing
  - Machine learning scenarios such as Spark in-memory computing and MLlib
  - Search and log data processing scenarios in which solutions such as Elasticsearch and Kafka are used
- This instance family supports online replacement and hot swapping of damaged disks to prevent instance shutdown.

  If a local disk fails, you receive a system event. You can handle the system event by initiating the process of repairing the damaged disk. For more information, see O&M scenarios and system events for instances equipped with local disks.

  > **Important**
  >
  > After you initiate the process of repairing the damaged disk, data stored on the damaged disk cannot be restored.

- Compute:
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors.
- Storage:
  - Is an instance family in which all instances are I/O optimized.
  - Supports enhanced SSDs (ESSDs), ESSD AutoPL disks, standard SSDs, and ultra disks.
- Network:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

d2c instance types

Instance type	vCPUs	Memory size (GiB)	Local storage (GB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)
ecs.d2c.6xlarge	24	88.0	3 * 3,972	12.0	1,600,000
ecs.d2c.12xlarge	48	176.0	6 * 3,972	20.0	2,000,000
ecs.d2c.24xlarge	96	352.0	12 * 3,972	35.0	4,500,000

d2s, storage-intensive big data instance family

Features:

- This instance family is equipped with high-capacity and high-throughput local SATA HDDs and can provide a maximum bandwidth of 35 Gbit/s between instances.
- Supported scenarios:
  - Big data computing and storage business scenarios in which services such as Hadoop MapReduce, HDFS, Hive, and HBase are used
  - Machine learning scenarios such as Spark in-memory computing and MLlib
  - Search and log data processing scenarios in which solutions such as Elasticsearch and Kafka are used
- This instance family supports online replacement and hot swapping of damaged disks to prevent instance shutdown.

  If a local disk fails, you receive a system event. You can handle the system event by initiating the process of repairing the damaged disk. For more information, see O&M scenarios and system events for instances equipped with local disks.

  > **Important**
  >
  > After you initiate the process of repairing the damaged disk, data stored on the damaged disk cannot be restored.

- Compute:
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors.
- Storage:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks.
- Network:

- Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
- Provides high network performance based on large computing capacity.

d2s instance types

Instance type	vCPUs	Memory size (GiB)	Local storage (GB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)
ecs.d2s.5xlarge	20	88.0	8 * 7,838	12.0	1,600,000
ecs.d2s.10xlarge	40	176.0	15 * 7,838	20.0	2,000,000
ecs.d2s.20xlarge	80	352.0	30 * 7,838	35.0	4,500,000

d1ne, network-enhanced big data instance family

Features:

- This instance family is equipped with high-capacity and high-throughput local SATA HDDs and can provide a maximum bandwidth of 35 Gbit/s between instances.
- Supported scenarios:
  - Scenarios in which services such as Hadoop MapReduce, HDFS, Hive, and HBase are used
  - Machine learning scenarios such as Spark in-memory computing and MLlib
  - Search and log data processing scenarios in which solutions such as Elasticsearch are used
- Compute:
  - Offers a CPU-to-memory ratio of 1:4, which is designed for big data scenarios.
  - Uses 2.5 GHz Intel® Xeon® E5-2682 v4 (Broadwell) processors.
- Storage:
  - Is an instance family in which all instances are I/O optimized.
  - Supports only standard SSDs and ultra disks.
- Network:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

d1ne instance types

Instance type	vCPUs	Memory size (GiB)	Local storage (GB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)
ecs.d1ne.2xlarge	8	32.0	4 * 5,905	6.0	1,000,000
ecs.d1ne.4xlarge	16	64.0	8 * 5,905	12.0	1,600,000
ecs.d1ne.6xlarge	24	96.0	12 * 5,905	16.0	2,000,000

ecs.d1ne-c8d3.8xlarge	32	128.0	12 * 5,905	20.0	2,000,000
ecs.d1ne.8xlarge	32	128.0	16 * 5,905	20.0	2,500,000
ecs.d1ne-c14d3.14xlarge	56	160.0	12 * 5,905	35.0	4,500,000
ecs.d1ne.14xlarge	56	224.0	28 * 5,905	35.0	4,500,000

d1, big data instance family

Features:

- This instance family is equipped with high-capacity and high-throughput local SATA HDDs and can provide a maximum bandwidth of 17 Gbit/s between instances.
- Supported scenarios:
  - Scenarios in which services such as Hadoop MapReduce, HDFS, Hive, and HBase are used
  - Machine learning scenarios such as Spark in-memory computing and MLlib
  - Scenarios in which customers in industries such as Internet and finance need to compute, store, and analyze big data
  - Search and log data processing scenarios in which solutions such as Elasticsearch are used
- Compute:
  - Offers a CPU-to-memory ratio of 1:4, which is designed for big data scenarios.
  - Uses 2.5 GHz Intel® Xeon® E5-2682 v4 (Broadwell) processors.
- Storage:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks.
- Network:
  - Supports IPv4
  - Provides high network performance based on large computing capacity.

d1 instance types

Instance type	vCPUs	Memory size (GiB)	Local storage (GB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)
ecs.d1.2xlarge	8	32.0	4 * 5,905	3.0	300,000
ecs.d1.3xlarge	12	48.0	6 * 5,905	4.0	400,000
ecs.d1.4xlarge	16	64.0	8 * 5,905	6.0	600,000
ecs.d1.6xlarge	24	96.0	12 * 5,905	8.0	800,000

ecs.d1-c8d3.8xlarge	32	128.0	12 * 5,905	10.0	1,000,000
ecs.d1.8xlarge	32	128.0	16 * 5,905	10.0	1,000,000
ecs.d1-c14d3.14xlarge	56	160.0	12 * 5,905	17.0	1,800,000
ecs.d1.14xlarge	56	224.0	28 * 5,905	17.0	1,800,000

i4, instance family with local SSDs

- **Introduction**: This instance family is equipped with high-performance local NVMe SSDs that deliver high IOPS, high I/O throughput, and low latency.
- **Supported scenarios**: OLTP and high-performance relational databases, NoSQL databases such as Cassandra and MongoDB, and search scenarios that use solutions such as Elasticsearch.
- **Compute**:
  - Uses 2.7 GHz Intel® Xeon® Scalable (Ice Lake) processors that deliver an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.
- **Is compatible with specific operating systems**. For more information, see Compatibility between the i4 instance types and operating systems.

i4 instance types

Instance type	vCPUs	Memory (GiB)	Local storage (GB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.i4.large	2	16	1 * 479	2.5/15	900,000	20,000/burstable up to 110,000	1.5/6
ecs.i4.xlarge	4	32	1 * 959	4/15	1,000,000	40,000/burstable up to 110,000	2/6
ecs.i4.2xlarge	8	64	1 * 1919	6/15	1,600,000	50,000/burstable up to 110,000	3/6
ecs.i4.4xlarge	16	128	1 * 3837	10/25	3,000,000	80,000/burstable up to 110,000	5/6

ecs.i4.8xlarge	32	256	2 * 3837	25/none	6,000,000	150,000/none	8/none
ecs.i4.16xlarge	64	512	4 * 3837	50/none	12,000,000	300,000/none	16/none
ecs.i4.32xlarge	128	1024	8 * 3837	100/none	24,000,000	600,000/none	32/none

i4g, instance family with local SSDs

- **Introduction**: This instance family is equipped with high-performance local NVMe SSDs that deliver high IOPS, high I/O throughput, and low latency.
- **Supported scenarios**: OLTP and high-performance relational databases, E-MapReduce big data scenarios such as tiering of hot and cold data, storage and computing separation, and data lakes, and search scenarios that use solutions such as Elasticsearch.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4, which is designed for high-performance databases.
  - Uses 2.7 GHz Intel® Xeon® Scalable (Ice Lake) processors that deliver an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.

i4g instance types

Instance type	vCPUs	Memory (GiB)	Local storage (GB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.i4g.4xlarge	16	64	1 * 959	8/25	3,000,000	100,000	6
ecs.i4g.8xlarge	32	128	1 * 1919	16/25	6,000,000	150,000	8
ecs.i4g.16xlarge	64	256	2 * 1919	32/none	12,000,000	300,000	16
ecs.i4g.32xlarge	128	512	4 * 1919	64/none	24,000,000	600,000	32

> **Note**
>
> This instance family supports only Linux images. When you create an instance of this instance family, you must select a Linux image.

i4r, instance family with local SSDs

- **Introduction**: This instance family is equipped with high-performance local NVMe SSDs that deliver high IOPS, high I/O throughput, and low latency.
- **Supported scenarios**: OLTP and high-performance relational databases, NoSQL databases such as Cassandra and MongoDB, and search scenarios that use solutions such as Elasticsearch.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8, which is designed for high-performance databases. This instance family is the most cost-effective instance family that is suitable for scenarios such as hot data tiering and data lakes.
  - Uses 2.7 GHz Intel® Xeon® Scalable (Ice Lake) processors that deliver an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.

i4r instance types

Instance type	vCPUs	Memory (GiB)	Local storage (GB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.i4r.4xlarge	16	128	1 * 959	8/25	3,000,000	100,000	6
ecs.i4r.8xlarge	32	256	1 * 1919	16/25	6,000,000	150,000	8
ecs.i4r.16xlarge	64	512	2 * 1919	32/none	12,000,000	300,000	16
ecs.i4r.32xlarge	128	1024	4 * 1919	64/none	24,000,000	600,000	32

i4p, performance-enhanced instance family with local SSDs

- **Introduction**: This instance family uses the Intel® Second-generation Optane persistent memory (BPS) to provide ultra-high-performance local disks. For information about how to initialize local disks, see the Configure persistent memory as a local disk section of the "Configure the usage mode of persistent memory" topic.
- **Supported scenarios**:
  - Gene sequencing applications. For more information, see Case description.
  - On-disk key-value (KV) databases, such as RocksDB and ClickHouse.
  - OLTP and high-performance relational databases for write-ahead log (WAL) optimization.
  - NoSQL databases, such as Cassandra, MongoDB, and HBase.
  - Search scenarios that use solutions such as Elasticsearch.

- Other I/O-intensive applications that frequently write data to disks, such as message middleware and containers.

- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses the third-generation Intel® Xeon® Scalable (Ice Lake) processors that deliver a base frequency of 2.7 GHz and an all-core turbo frequency of 3.2 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

i4p instance types

Instance type	vCPUs	Memory (GiB)	Persistent memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.i4p.2xlarge	8	32	1 * 126	5/10	1,600,000	50,000/burstable up to 110,000	3/6
ecs.i4p.4xlarge	16	64	2 * 126	10/25	3,000,000	80,000/burstable up to 110,000	5/6
ecs.i4p.6xlarge	24	96	3 * 126	12/25	4,500,000	110,000/none	6/none
ecs.i4p.8xlarge	32	128	4 * 126	16/25	6,000,000	150,000/none	8/none
ecs.i4p.16xlarge	64	256	1 * 1008	32/none	12,000,000	300,000/none	16/none
ecs.i4p.32xlarge	128	512	2 * 1008	64/none	24,000,000	600,000/none	32/none

i3g, instance family with local SSDs

- **Introduction**: This instance family is equipped with high-performance local NVMe SSDs that deliver high IOPS, high I/O throughput, and low latency.
- **Supported scenarios**: OLTP and high-performance relational databases, NoSQL databases such as Cassandra, MongoDB, and HBase, and search scenarios that use solutions such as Elasticsearch.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4, which is designed for high-performance databases.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver a turbo frequency of 3.2 GHz to provide consistent computing performance.
- **Storage**:

- Is an instance family in which all instances are I/O optimized.
- Supports ESSDs and ESSD AutoPL disks.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

i3g instance types

Instance type	vCPUs	Memory (GiB)	Local storage (GB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.i3g.2xlarge	8	32	1 * 479	3/10	1,750,000	52,500	2
ecs.i3g.4xlarge	16	64	1 * 959	5/10	3,500,000	84,000	3
ecs.i3g.8xlarge	32	128	2 * 959	12/none	7,000,000	157,500	5
ecs.i3g.13xlarge	52	192	3 * 959	16/none	12,000,000	252,000	8
ecs.i3g.26xlarge	104	384	6 * 959	32/none	24,000,000	500,000	16

Note

This instance family supports only Linux images. When you create an instance of this instance family, you must select a Linux image.

i3, instance family with local SSDs

- **Introduction**: This instance family is equipped with high-performance local NVMe SSDs that deliver high IOPS, high I/O throughput, and low latency, and allows damaged disks to be isolated online.
- **Supported scenarios**: OLTP and high-performance relational databases, NoSQL databases such as Cassandra and MongoDB, and search scenarios that use solutions such as Elasticsearch.
- **Compute**:
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver a turbo frequency of 3.2 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

i3 instance types

Instance type	vCPUs	Memory (GiB)	Local storage (GB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.i3.xlarge	4	32	1 * 959	1.5/10	1,000,000	40,000	1.5
ecs.i3.2xlarge	8	64	1 * 1919	2.5/10	1,600,000	50,000	2
ecs.i3.4xlarge	16	128	2 * 1919	5/10	3,000,000	80,000	3
ecs.i3.8xlarge	32	256	4 * 1919	10/none	6,000,000	150,000	5
ecs.i3.13xlarge	52	384	6 * 1919	16/none	9,000,000	240,000	8
ecs.i3.26xlarge	104	768	12 * 1919	32/none	24,000,000	480,000	16

Note

This instance family supports only Linux images. When you create an instance of this instance family, you must select a Linux image.

i2, instance family with local SSDs

- **Introduction**: This instance family is equipped with high-performance local NVMe SSDs that deliver high IOPS, high I/O throughput, and low latency.
- **Supported scenarios**: OLTP and high-performance relational databases, NoSQL databases such as Cassandra, MongoDB, and HBase, and search scenarios that use solutions such as Elasticsearch.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8, which is designed for high-performance databases.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

i2 instance types

Instance type	vCPUs	Memory (GiB)	Local storage (GB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Disk bandwidth (Gbit/s)
ecs.i2.xlarge	4	32	1 * 959	1	500,000	Up to 16
ecs.i2.2xlarge	8	64	1 * 1919	2	1,000,000	Up to 16
ecs.i2.4xlarge	16	128	2 * 1919	3	1,500,000	Up to 16
ecs.i2.8xlarge	32	256	4 * 1919	6	2,000,000	Up to 16
ecs.i2.16xlarge	64	512	8 * 1919	10	4,000,000	Up to 16

i2g, instance family with local SSDs

- **Introduction**: This instance family is equipped with high-performance local NVMe SSDs that deliver high IOPS, high I/O throughput, and low latency.
- **Supported scenarios**: OLTP and high-performance relational databases, NoSQL databases such as Cassandra, MongoDB, and HBase, and search scenarios that use solutions such as Elasticsearch.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4, which is designed for high-performance databases.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks.
- **Network**:
  - Supports only IPv4.
  - Provides high network performance based on large computing capacity.

i2g instance types

Instance type	vCPUs	Memory (GiB)	Local storage (GB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)
ecs.i2g.2xlarge	8	32	1 * 959	2	1,000,000
ecs.i2g.4xlarge	16	64	1 * 1919	3	1,500,000
ecs.i2g.8xlarge	32	128	2 * 1919	6	2,000,000
ecs.i2g.16xlarge	64	256	4 * 1919	10	4,000,000

i2ne, instance family with local SSDs

- **Introduction**: This instance family is equipped with high-performance local NVMe SSDs that deliver high IOPS, high I/O throughput, and low latency.
- **Supported scenarios**: OLTP and high-performance relational databases, NoSQL databases such as Cassandra, MongoDB, and HBase, and search scenarios that use solutions such as Elasticsearch.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8, which is designed for high-performance databases.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.
  - Provides a network bandwidth of up to 20 Gbit/s.

i2ne instance types

Instance type	vCPUs	Memory (GiB)	Local storage (GB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Disk bandwidth (Gbit/s)
ecs.i2ne.xlarge	4	32	1 * 959	1.5	500,000	Up to 16
ecs.i2ne.2xlarge	8	64	1 * 1919	2.5	1,000,000	Up to 16
ecs.i2ne.4xlarge	16	128	2 * 1919	5	1,500,000	Up to 16
ecs.i2ne.8xlarge	32	256	4 * 1919	10	2,000,000	Up to 16
ecs.i2ne.16xlarge	64	512	8 * 1919	20	4,000,000	Up to 16
ecs.i2ne.20xlarge	80	704	10 * 1919	25	4,500,000	Up to 16

i2gne, instance family with local SSDs

- **Introduction**: This instance family is equipped with high-performance local NVMe SSDs that deliver high IOPS, high I/O throughput, and low latency.
- **Supported scenarios**: OLTP and high-performance relational databases, NoSQL databases such as Cassandra, MongoDB, and HBase, and search scenarios that use solutions such as Elasticsearch.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4, which is designed for high-performance databases.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.
  - Provides a network bandwidth of up to 20 Gbit/s.

i2gne instance types

Instance type	vCPUs	Memory (GiB)	Local storage (GB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)
ecs.i2gne.2xlarge	8	32	1 * 959	2.5	1,000,000
ecs.i2gne.4xlarge	16	64	1 * 1919	5	1,500,000
ecs.i2gne.8xlarge	32	128	2 * 1919	10	2,000,000
ecs.i2gne.16xlarge	64	256	4 * 1919	20	4,000,000

i1, instance family with local SSDs

- **Introduction**: This instance family is equipped with high-performance local NVMe SSDs that deliver high IOPS, high I/O throughput, and low latency.
- **Supported scenarios**: OLTP and high-performance relational databases, NoSQL databases such as Cassandra and MongoDB, and search scenarios that use solutions such as Elasticsearch.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4, which is designed for high-performance databases.
  - Uses 2.5 GHz Intel® Xeon® E5-2682 v4 (Broadwell) processors.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks.
- **Network**:
  - Supports only IPv4.
  - Provides high network performance based on large computing capacity.

i1 instance types

Instance type	vCPUs	Memory (GiB)	Local storage (GB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)
ecs.i1.xlarge	4	16	2 * 111	0.8	200,000

ecs.i1.2xlarge	8	32	2 * 223	1.5	400,000
ecs.i1.3xlarge	12	48	2 * 335	2	400,000
ecs.i1.4xlarge	16	64	2 * 446	3	500,000
ecs.i1-c5d1.4xlarge	16	64	2 * 1563	3	400,000
ecs.i1.6xlarge	24	96	2 * 670	4.5	600,000
ecs.i1.8xlarge	32	128	2 * 893	6	800,000
ecs.i1-c10d1.8xlarge	32	128	2 * 1563	6	800,000
ecs.i1.14xlarge	56	224	2 * 1563	10	1,200,000

hfc7, compute-optimized instance family with high clock speeds

- **Introduction**: This instance family offloads a large number of virtualization features to dedicated hardware by using the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads.
- **Supported scenarios**:
  - Scenarios in which large volumes of packets are received and transmitted, such as live commenting and telecom data forwarding
  - High-performance frontend server clusters
  - Frontend servers for MMO games
  - Data analytics, batch processing, and video encoding
  - High-performance scientific and engineering applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses Intel® Xeon® Cooper Lake processors that deliver an all-core turbo frequency of 3.8 GHz and have a minimum clock speed of 3.3 GHz to provide consistent computing performance.
  - Allows you to enable or disable Hyper-Threading.

    > **Note**
    >
    > By default, Hyper-Threading is enabled for ECS instances. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports only ESSDs and ESSD AutoPL disks.

- Provides high storage I/O performance based on large computing capacity.

> **Note**
>
> For information about the storage I/O performance of the next-generation, enterprise-level instance families, see Storage I/O performance.

- **Network**:
  - Supports IPv4 and IPv6.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The hfc7 instance family includes the following instance types: ecs.hfc7.large, ecs.hfc7.xlarge, ecs.hfc7.2xlarge, ecs.hfc7.3xlarge, ecs.hfc7.4xlarge, ecs.hfc7.6xlarge, ecs.hfc7.8xlarge, ecs.hfc7.12xlarge, and ecs.hfc7.24xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.hfc7.large	2	4	1.2/10	900,000	250,000	2	2	6	6	20,000	1
ecs.hfc7.xlarge	4	8	2/10	1,000,000	250,000	4	3	15	15	30,000	1.5
ecs.hfc7.2xlarge	8	16	3/10	1,600,000	250,000	8	4	15	15	45,000	2
ecs.hfc7.3xlarge	12	24	4.5/10	2,000,000	250,000	8	6	15	15	60,000	2.5
ecs.hfc7.4xlarge	16	32	6/10	2,500,000	300,000	8	8	30	30	75,000	3
ecs.hfc7.6xlarge	24	48	8/10	3,000,000	450,000	12	8	30	30	90,000	4
ecs.hfc7.8xlarge	32	64	10/none	4,000,000	600,000	16	8	30	30	105,000	5

ecs.hfc7.12xlarge	48	96	16/none	6,000,000	1,000,000	24	8	30	30	150,000	8
ecs.hfc7.24xlarge	96	192	32/none	12,000,000	1,800,000	32	15	30	30	300,000	16

hfc6, compute-optimized instance family with high clock speeds

- **Introduction**: This instance family offloads a large number of virtualization features to dedicated hardware by using the SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads.
- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Web frontend servers
  - Frontend servers for MMO games
  - Data analytics, batch processing, and video encoding
  - High-performance scientific and engineering applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 3.1 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver a turbo frequency of 3.5 GHz to provide consistent computing performance.

    > **Note**
    >
    > The processors used by this instance family have a clock speed of 3.1 GHz. However, the Intel System Studio (ISS) feature may cause a lower clock speed to be displayed. Alibaba Cloud is working on this issue. This issue does not affect the actual clock speeds of your instances.
    >
    > You can separately run the following commands to use the turbostat tool to view the actual clock speeds:
    >
    > ```
    > yum install kernel-tools
    > ```
    >
    > ```
    > turbostat
    > ```

  - Allows you to enable or disable Hyper-Threading.

    > **Note**
    >
    > By default, Hyper-Threading is enabled for ECS instances. For more information, see Specify and view CPU options.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs), ESSD AutoPL disks, standard SSDs, and ultra disks.

- Provides high storage I/O performance based on large computing capacity.

> **Note**
>
> For information about the storage I/O performance of the next-generation, enterprise-level instance families, see Storage I/O performance.

- **Network**:
  - Supports IPv4 and IPv6.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The hfc6 instance family includes the following instance types: ecs.hfc6.large, ecs.hfc6.xlarge, ecs.hfc6.2xlarge, ecs.hfc6.3xlarge, ecs.hfc6.4xlarge, ecs.hfc6.6xlarge, ecs.hfc6.8xlarge, ecs.hfc6.10xlarge, ecs.hfc6.16xlarge, and ecs.hfc6.20xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.hfc6.large	2	4	1/3	300,000	35,000	2	2	6	1	10,000	1
ecs.hfc6.xlarge	4	8	1.5/5	500,000	70,000	4	3	10	1	20,000	1.5
ecs.hfc6.2xlarge	8	16	2.5/8	800,000	150,000	8	4	10	1	25,000	2
ecs.hfc6.3xlarge	12	24	4/10	900,000	220,000	8	6	10	1	30,000	2.5
ecs.hfc6.4xlarge	16	32	5/10	1,000,000	300,000	8	8	20	1	40,000	3
ecs.hfc6.6xlarge	24	48	7.5/10	1,500,000	450,000	12	8	20	1	50,000	4
ecs.hfc6.8xlarge	32	64	10/none	2,000,000	600,000	16	8	20	1	60,000	5

ecs.hfc6.10xlarge	40	96	12.5/none	3,000,000	1,000,000	32	7	20	1	100,000	8
ecs.hfc6.16xlarge	64	128	20/none	4,000,000	1,200,000	32	8	20	1	120,000	10
ecs.hfc6.20xlarge	80	192	25/none	6,000,000	1,800,000	32	15	20	1	200,000	16

hfg7, general-purpose instance family with high clock speeds

- **Introduction**: This instance family offloads a large number of virtualization features to dedicated hardware by using the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads.
- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting and telecom data forwarding
  - Enterprise-level applications of various types and sizes
  - Game servers
  - Small and medium-sized database systems, caches, and search clusters
  - High-performance scientific computing
  - Video encoding applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses Intel® Xeon® Cooper Lake processors that deliver an all-core turbo frequency of 3.8 GHz and have a minimum clock speed of 3.3 GHz to provide consistent computing performance.
  - Allows you to enable or disable Hyper-Threading.

    > **Note**
    >
    > By default, Hyper-Threading is enabled for ECS instances. For more information, see Specify and view CPU options.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports only ESSDs and ESSD AutoPL disks.
  - Provides high storage I/O performance based on large computing capacity.

    > **Note**
    >
    > For information about the storage I/O performance of the next-generation, enterprise-level instance families, see Storage I/O performance.

- **Network**:
  - Supports IPv4 and IPv6.
  - Provides ultra-high packet forwarding rates.

- Provides high network performance based on large computing capacity.

The hfg7 instance family includes the following instance types: ecs.hfg7.large, ecs.hfg7.xlarge, ecs.hfg7.2xlarge, ecs.hfg7.3xlarge, ecs.hfg7.4xlarge, ecs.hfg7.6xlarge, ecs.hfg7.8xlarge, ecs.hfg7.12xlarge, and ecs.hfg7.24xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.hfg7.large	2	8	1.2/10	900,000	250,000	2	2	6	6	20,000	1
ecs.hfg7.xlarge	4	16	2/10	1,000,000	250,000	4	3	15	15	30,000	1.5
ecs.hfg7.2xlarge	8	32	3/10	1,600,000	250,000	8	4	15	15	45,000	2
ecs.hfg7.3xlarge	12	48	4.5/10	2,000,000	250,000	8	6	15	15	60,000	2.5
ecs.hfg7.4xlarge	16	64	6/10	2,500,000	300,000	8	8	30	30	75,000	3
ecs.hfg7.6xlarge	24	96	8/10	3,000,000	450,000	12	8	30	30	90,000	4
ecs.hfg7.8xlarge	32	128	10/none	4,000,000	600,000	16	8	30	30	105,000	5
ecs.hfg7.12xlarge	48	192	16/none	6,000,000	1,000,000	24	8	30	30	150,000	8
ecs.hfg7.24xlarge	96	384	32/none	12,000,000	1,800,000	32	15	30	30	300,000	16

hfg6, general-purpose instance family with high clock speeds

- **Introduction**: This instance family offloads a large number of virtualization features to dedicated hardware by using the SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads.
- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Enterprise-level applications of various types and sizes
  - Websites and application servers
  - Game servers
  - Small and medium-sized database systems, caches, and search clusters
  - Data analytics and computing
  - Computing clusters and memory-intensive data processing
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 3.1 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver a turbo frequency of 3.5 GHz to provide consistent computing performance.

    > **Note**
    >
    > The processors used by this instance family have a clock speed of 3.1 GHz. However, the Intel System Studio (ISS) feature may cause a lower clock speed to be displayed. Alibaba Cloud is working on this issue. This issue does not affect the actual clock speeds of your instances.
    >
    > You can separately run the following commands to use the turbostat tool to view the actual clock speeds:
    >
    > ```
    > yum install kernel-tools
    > ```
    >
    > ```
    > turbostat
    > ```

  - Allows you to enable or disable Hyper-Threading.

    > **Note**
    >
    > By default, Hyper-Threading is enabled for ECS instances. For more information, see Specify and view CPU options.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks.
  - Provides high storage I/O performance based on large computing capacity.

    > **Note**
    >
    > For information about the storage I/O performance of the next-generation, enterprise-level instance families, see Storage I/O performance.

- **Network**:
  - Supports IPv4 and IPv6.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The hfg6 instance family includes the following instance types: ecs.hfg6.large, ecs.hfg6.xlarge, ecs.hfg6.2xlarge, ecs.hfg6.3xlarge, ecs.hfg6.4xlarge, ecs.hfg6.6xlarge, ecs.hfg6.8xlarge, ecs.hfg6.10xlarge, ecs.hfg6.16xlarge, and ecs.hfg6.20xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance type

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.hfg6.large	2	8	1/3	300,000	35,000	2	2	6	1	10,000	1
ecs.hfg6.xlarge	4	16	1.5/5	500,000	70,000	4	3	10	1	20,000	1.5
ecs.hfg6.2xlarge	8	32	2.5/8	800,000	150,000	8	4	10	1	25,000	2
ecs.hfg6.3xlarge	12	48	4/10	900,000	220,000	8	6	10	1	30,000	2.5
ecs.hfg6.4xlarge	16	64	5/10	1,000,000	300,000	8	8	20	1	40,000	3
ecs.hfg6.6xlarge	24	96	7.5/10	1,500,000	450,000	12	8	20	1	50,000	4
ecs.hfg6.8xlarge	32	128	10/none	2,000,000	600,000	16	8	20	1	60,000	5
ecs.hfg6.10xlarge	40	192	12.5/none	3,000,000	1,000,000	32	7	20	1	100,000	8
ecs.hfg6.16xlarge	64	256	20/none	4,000,000	1,200,000	32	8	20	1	120,000	10
ecs.hfg6.20xlarge	80	384	25/none	6,000,000	1,800,000	32	15	20	1	200,000	16

hfr7, memory-optimized instance family with high clock speeds

- **Introduction**: This instance family offloads a large number of virtualization features to dedicated hardware by using the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads.
- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - High-performance databases and in-memory databases
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses Intel® Xeon® Cooper Lake processors that deliver an all-core turbo frequency of 3.8 GHz and have a minimum clock speed of 3.3 GHz to provide consistent computing performance.
  - Allows you to enable or disable Hyper-Threading.

    > **Note**
    >
    > By default, Hyper-Threading is enabled for ECS instances. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports only ESSDs and ESSD AutoPL disks.
  - Provides high storage I/O performance based on large computing capacity.

    > **Note**
    >
    > For information about the storage I/O performance of the next-generation, enterprise-level instance families, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The hfr7 instance family includes the following instance types: ecs.hfr7.large, ecs.hfr7.xlarge, ecs.hfr7.2xlarge, ecs.hfr7.3xlarge, ecs.hfr7.4xlarge, ecs.hfr7.6xlarge, ecs.hfr7.8xlarge, ecs.hfr7.12xlarge, and ecs.hfr7.24xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)

ecs.hfr7.large	2	16	1.2/10	900,000	250,000	2	2	6	6	20,000	1
ecs.hfr7.xlarge	4	32	2/10	1,000,000	250,000	4	3	15	15	30,000	1.5
ecs.hfr7.2xlarge	8	64	3/10	1,600,000	250,000	8	4	15	15	45,000	2
ecs.hfr7.3xlarge	12	96	4.5/10	2,000,000	250,000	8	6	15	15	60,000	2.5
ecs.hfr7.4xlarge	16	128	6/10	2,500,000	300,000	8	8	30	30	75,000	3
ecs.hfr7.6xlarge	24	192	8/10	3,000,000	450,000	12	8	30	30	90,000	4
ecs.hfr7.8xlarge	32	256	10/none	4,000,000	600,000	16	8	30	30	105,000	5
ecs.hfr7.12xlarge	48	384	16/none	6,000,000	1,000,000	24	8	30	30	150,000	8
ecs.hfr7.24xlarge	96	768	32/none	12,000,000	1,800,000	32	15	30	30	300,000	16

hfr6, memory-optimized instance family with high clock speeds

- **Introduction**: This instance family offloads a large number of virtualization features to dedicated hardware by using the SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads.
- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - High-performance databases and in-memory databases
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.

- Uses 3.1 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver a turbo frequency of 3.5 GHz to provide consistent computing performance.

  > **Note**
  >
  > The processors used by this instance family have a clock speed of 3.1 GHz. However, the Intel System Studio (ISS) feature may cause a lower clock speed to be displayed. Alibaba Cloud is working on this issue. This issue does not affect the actual clock speeds of your instances.
  >
  > You can separately run the following commands to use the turbostat tool to view the actual clock speeds:
  >
  > ```
  > yum install kernel-tools
  > ```
  >
  > ```
  > turbostat
  > ```

- Allows you to enable or disable Hyper-Threading.

  > **Note**
  >
  > By default, Hyper-Threading is enabled for ECS instances. For more information, see Specify and view CPU options.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks.
  - Provides high storage I/O performance based on large computing capacity.

    > **Note**
    >
    > For information about the storage I/O performance of the next-generation, enterprise-level instance families, see Storage I/O performance.

- **Network**:
  - Supports IPv4 and IPv6.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The hfr6 instance family includes the following instance types: ecs.hfr6.large, ecs.hfr6.xlarge, ecs.hfr6.2xlarge, ecs.hfr6.3xlarge, ecs.hfr6.4xlarge, ecs.hfr6.6xlarge, ecs.hfr6.8xlarge, ecs.hfr6.10xlarge, ecs.hfr6.16xlarge, and ecs.hfr6.20xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.hfr6.large	2	16	1/3	300,000	35,000	2	2	6	1	10,000	1

ecs.hfr6.xlarge	4	32	1.5/5	500,000	70,000	4	3	10	1	20,000	1.5
ecs.hfr6.2xlarge	8	64	2.5/8	800,000	150,000	8	4	10	1	25,000	2
ecs.hfr6.3xlarge	12	96	4/10	900,000	220,000	8	6	10	1	30,000	2.5
ecs.hfr6.4xlarge	16	128	5/10	1,000,000	300,000	8	8	20	1	40,000	3
ecs.hfr6.6xlarge	24	192	7.5/10	1,500,000	450,000	12	8	20	1	50,000	4
ecs.hfr6.8xlarge	32	256	10/none	2,000,000	600,000	16	8	20	1	60,000	5
ecs.hfr6.10xlarge	40	384	12.5/none	3,000,000	1,000,000	32	7	20	1	100,000	8
ecs.hfr6.16xlarge	64	512	20/none	4,000,000	1,200,000	32	8	20	1	120,000	10
ecs.hfr6.20xlarge	80	768	25/none	6,000,000	1,800,000	32	15	20	1	200,000	16

hfc5, compute-optimized instance family with high clock speeds

- **Supported scenarios**: Scenarios such as high-performance frontend servers, high-performance scientific and engineering applications, MMO games, and video encoding.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 3.1 GHz Intel® Xeon® Gold 6149 (Skylake) processors.
  - Provides consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks.
- **Network**:
  - Supports only IPv4.
  - Provides high network performance based on large computing capacity.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI
ecs.hfc5.large	2	4	1	300,000	2	2	6
ecs.hfc5.xlarge	4	8	1.5	500,000	2	3	10
ecs.hfc5.2xlarge	8	16	2	1,000,000	2	4	10
ecs.hfc5.3xlarge	12	24	2.5	1,300,000	4	6	10
ecs.hfc5.4xlarge	16	32	3	1,600,000	4	8	20
ecs.hfc5.6xlarge	24	48	4.5	2,000,000	6	8	20
ecs.hfc5.8xlarge	32	64	6	2,500,000	8	8	20

hfg5, general-purpose instance family with high clock speeds

- **Supported scenarios**: Scenarios such as high-performance frontend servers, high-performance scientific and engineering applications, MMO games, and video encoding.

Compute:

- Offers a CPU-to-memory ratio of 1:4 (excluding the instance type with 56 vCPUs).
- Uses 3.1 GHz Intel® Xeon® Gold 6149 (Skylake) processors.
- Provides consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks.
- **Network**:
  - Supports only IPv4.
  - Provides high network performance based on large computing capacity.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI
ecs.hfg5.large	2	8	1	300,000	2	2	6
ecs.hfg5.xlarge	4	16	1.5	500,000	2	3	10
ecs.hfg5.2xlarge	8	32	2	1,000,000	2	4	10
ecs.hfg5.3xlarge	12	48	2.5	1,300,000	4	6	10
ecs.hfg5.4xlarge	16	64	3	1,600,000	4	8	20
ecs.hfg5.6xlarge	24	96	4.5	2,000,000	6	8	20
ecs.hfg5.8xlarge	32	128	6	2,500,000	8	8	20
ecs.hfg5.14xlarge	56	160	10	4,000,000	14	8	20

g7se, storage-enhanced general-purpose instance family

- **Introduction**: This instance family uses the third-generation SHENLONG architecture and Intel Ice Lake processors to improve storage I/O performance.
- **Supported scenarios**: I/O-intensive scenarios such as large and medium-sized online transaction processing (OLTP) core databases, large and medium-sized NoSQL databases, search and real-time log analytics, and traditional large enterprise-level commercial software such as SAP.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses the third-generation Intel® Xeon® Scalable (Ice Lake) processors that deliver a base frequency of 2.9 GHz and an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.

- Allows up to 64 data disks to be attached per instance. You can attach up to 16 data disks to an instance when you create the instance. If the instance requires more data disks, attach more data disks after the instance is created. For more information, see Attach a data disk.
- Delivers a sequential read/write throughput of up to 64 Gbit/s and IOPS of up to 1,000,000 per instance.
- Provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The g7se instance family includes the following instance types: ecs.g7se.large, ecs.g7se.xlarge, ecs.g7se.2xlarge, ecs.g7se.3xlarge, ecs.g7se.4xlarge, ecs.g7se.6xlarge, ecs.g7se.8xlarge, and ecs.g7se.16xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Maximum attached data disks	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.g7se.large	2	8	1.2/burstable up to 3	450,000	Up to 250,000	2	3	6	6	16	30,000/burstable up to 150,000	3/10
ecs.g7se.xlarge	4	16	2/burstable up to 5	500,000	Up to 250,000	4	4	15	15	16	60,000/burstable up to 150,000	4/10
ecs.g7se.2xlarge	8	32	3/burstable up to 8	800,000	Up to 250,000	8	4	15	15	16	100,000/burstable up to 150,000	6/10
ecs.g7se.3xlarge	12	48	4.5/burstable up to 10	1,200,000	Up to 250,000	8	8	15	15	16	120,000/burstable up to 150,000	8/10

ecs.g7se.4xlarge	16	64	6/burstable up to 10	1,500,000	300,000	8	8	30	30	24	150,000/none	10/none
ecs.g7se.6xlarge	24	96	8/burstable up to 10	2,250,000	450,000	12	8	30	30	24	200,000/none	12/none
ecs.g7se.8xlarge	32	128	10/none	3,000,000	600,000	16	8	30	30	30	300,000/none	16/none
ecs.g7se.16xlarge	64	256	16/none	6,000,000	1,200,000	32	8	30	30	56	500,000/none	32/none

c7se, storage-enhanced compute-optimized instance family

- **Introduction**: This instance family uses the third-generation SHENLONG architecture and Intel Ice Lake processors to improve storage I/O performance.
- **Supported scenarios**: I/O-intensive scenarios such as large and medium-sized OLTP core databases, large and medium-sized NoSQL databases, search and real-time log analytics, and traditional large enterprise-level commercial software such as SAP.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses the third-generation Intel® Xeon® Scalable (Ice Lake) processors that deliver a base frequency of 2.9 GHz and an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports the Non-Volatile Memory Express (NVMe) protocol. For more information, see NVMe protocol.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Allows up to 64 data disks to be attached per instance. You can attach up to 16 data disks to an instance when you create the instance. If the instance requires more data disks, attach more data disks after the instance is created. For more information, see Attach a data disk.
  - Delivers a sequential read/write throughput of up to 64 Gbit/s and IOPS of up to 1,000,000 per instance.
  - Provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The c7se instance family includes the following instance types: ecs.c7se.large, ecs.c7se.xlarge, ecs.c7se.2xlarge, ecs.c7se.3xlarge, ecs.c7se.4xlarge, ecs.c7se.6xlarge, ecs.c7se.8xlarge, and ecs.c7se.16xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Maximum attached data disks	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.c7se.large	2	4	1.2/burstable up to 3	450,000	Up to 250,000	2	3	6	6	16	30,000/burstable up to 150,000	3/10
ecs.c7se.xlarge	4	8	2/burstable up to 5	500,000	Up to 250,000	4	4	15	15	16	60,000/burstable up to 150,000	4/10
ecs.c7se.2xlarge	8	16	3/burstable up to 8	800,000	Up to 250,000	8	4	15	15	16	100,000/burstable up to 150,000	6/10
ecs.c7se.3xlarge	12	24	4.5/burstable up to 10	1,200,000	Up to 250,000	8	8	15	15	16	120,000/burstable up to 150,000	8/10
ecs.c7se.4xlarge	16	32	6/burstable up to 10	1,500,000	300,000	8	8	30	30	24	150,000/none	10/none
ecs.c7se.6xlarge	24	48	8/burstable up to 10	2,250,000	450,000	12	8	30	30	24	200,000/none	12/none
ecs.c7se.8xlarge	32	64	10/none	3,000,000	600,000	16	8	30	30	30	300,000/none	16/none

ecs.c7se.16xlarge	64	128	16/none	6,000,000	1,200,000	32	8	30	30	56	500,000/none	32/none

r7se, storage-enhanced memory-optimized instance family

- **Introduction**: This instance family uses the third-generation SHENLONG architecture and Intel Ice Lake processors to improve storage I/O performance.
- **Supported scenarios**:
  - I/O-intensive scenarios such as large and medium-sized OLTP core databases
  - Large and medium-sized NoSQL databases
  - Search and real-time log analytics
  - Traditional large enterprise-level commercial software such as SAP
  - High-density deployment of containers
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses the third-generation Intel® Xeon® Scalable (Ice Lake) processors that deliver a base frequency of 2.9 GHz and an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports the NVMe protocol. For more information, see NVMe protocol.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Allows up to 64 data disks to be attached per instance. You can attach up to 16 data disks to an instance when you create the instance. If the instance requires more data disks, attach more data disks after the instance is created. For more information, see Attach a data disk.
  - Delivers a sequential read/write throughput of up to 64 Gbit/s and IOPS of up to 1,000,000 per instance.
  - Provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The r7se instance family includes the following instance types: ecs.r7se.large, ecs.r7se.xlarge, ecs.r7se.2xlarge, ecs.r7se.3xlarge, ecs.r7se.4xlarge, ecs.r7se.6xlarge, ecs.r7se.8xlarge, and ecs.r7se.16xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Maximum attached data disks	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.r7se.large	2	16	1.2/burstable up to 3	450,000	Up to 250,000	2	3	6	6	16	30,000/burstable up to 150,000	3/10
ecs.r7se.xlarge	4	32	2/burstable up to 5	500,000	Up to 250,000	4	4	15	15	16	60,000/burstable up to 150,000	4/10
ecs.r7se.2xlarge	8	64	3/burstable up to 8	800,000	Up to 250,000	8	4	15	15	16	100,000/burstable up to 150,000	6/10
ecs.r7se.3xlarge	12	96	4.5/burstable up to 10	1,200,000	Up to 250,000	8	8	15	15	16	120,000/burstable up to 150,000	8/10
ecs.r7se.4xlarge	16	128	6/burstable up to 10	1,500,000	300,000	8	8	30	30	24	150,000/none	10/none
ecs.r7se.6xlarge	24	192	8/burstable up to 10	2,250,000	450,000	12	8	30	30	24	200,000/none	12/none

ecs.r7se.8xlarge	32	256	10/none	3,000,000	600,000	16	8	30	30	30	300,000/none	16/none
ecs.r7se.16xlarge	64	512	16/none	6,000,000	1,200,000	32	8	30	30	56	500,000/none	32/none

g7nex, network-enhanced general-purpose instance family

- **Introduction**: This instance family uses the fourth-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**:
  - Network-intensive scenarios such as Network Functions Virtualization (NFV) or Software-defined Wide Area Network (SD-WAN), mobile Internet, on-screen video comments, and telecom data forwarding
  - Small and medium-sized database systems, caches, and search clusters
  - Enterprise-level applications of various types and sizes
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses the third-generation Intel® Xeon® Scalable (Ice Lake) processors that deliver a base frequency of 2.7 GHz and an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Significantly improves the network throughput and packet forwarding rate per instance. A single instance can deliver a packet forwarding rate of up to 30,000,000 pps.
  - Provides high network performance based on large computing capacity.

The g7nex instance family includes the following instance types: ecs.g7nex.large, ecs.g7nex.xlarge, ecs.g7nex.2xlarge, ecs.g7nex.4xlarge, ecs.g7nex.8xlarge, ecs.g7nex.16xlarge, and ecs.g7nex.32xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	EBS queues	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.g7nex.large	2	8	3/burstable up to 20	450,000	2	3	10	10	1	10,000/burstable up to 50,000	1.5/burstable up to 8
ecs.g7nex.xlarge	4	16	5/burstable up to 24	900,000	4	4	15	15	1	20,000/burstable up to 50,000	2/burstable up to 8
ecs.g7nex.2xlarge	8	32	10/burstable up to 32	1,750,000	8	6	15	15	2	25,000/burstable up to 50,000	3/burstable up to 8
ecs.g7nex.4xlarge	16	64	20/burstable up to 40	3,000,000	16	8	30	30	2	40,000/burstable up to 50,000	5/burstable up to 8
ecs.g7nex.8xlarge	32	128	40/none	6,000,000	32	8	30	30	4	75,000/none	8/none
ecs.g7nex.16xlarge	64	256	80/none	8,000,000	32	15	50	50	4	150,000/none	16/none
ecs.g7nex.32xlarge	128	512	160/none	16,000,000	32	15	50	50	4	300,000/none	32/none

> **Note**
>
> Each ecs.g7nex.32xlarge instance must have at least two elastic network interfaces (ENIs) that are assigned different network card indexes before the instance can burst its network bandwidth to 160 Gbit/s. If all ENIs on the instance are assigned the same network card index, the instance can burst its network bandwidth only to 100 Gbit/s. For more information, see AttachNetworkInterface.

c7nex, network-enhanced compute-optimized instance family

- **Introduction**: This instance family uses the fourth-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**:
  - Network-intensive scenarios such as NFV or SD-WAN, mobile Internet, on-screen video comments, and telecom data forwarding
  - Small and medium-sized database systems, caches, and search clusters
  - Enterprise-level applications of various types and sizes
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses the third-generation Intel® Xeon® Scalable (Ice Lake) processors that deliver a base frequency of 2.7 GHz and an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- Network:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Significantly improves the network throughput and packet forwarding rate per instance. A single instance can deliver a packet forwarding rate of up to 30,000,000 pps.
  - Provides high network performance based on large computing capacity.
- **Security**: Supports the virtual Trusted Platform Module (vTPM) feature. For more information, see Overview of trusted computing capabilities.

The c7nex instance family includes the following instance types: ecs.c7nex.large, ecs.c7nex.xlarge, ecs.c7nex.2xlarge, ecs.c7nex.4xlarge, ecs.c7nex.8xlarge, ecs.c7nex.16xlarge, and ecs.c7nex.32xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	EBS queues	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.c7nex.large	2	4	3/burstable up to 20	450,000	2	3	10	10	1	10,000/burstable up to 50,000	1.5/burstable up to 8
ecs.c7nex.xlarge	4	8	5/burstable up to 24	900,000	4	4	15	15	1	20,000/burstable up to 50,000	2/burstable up to 8
ecs.c7nex.2xlarge	8	16	10/burstable up to 32	1,750,000	8	6	15	15	2	25,000/burstable up to 50,000	3/burstable up to 8
ecs.c7nex.4xlarge	16	32	20/burstable up to 40	3,000,000	16	8	30	30	2	40,000/burstable up to 50,000	5/burstable up to 8
ecs.c7nex.8xlarge	32	64	40/none	6,000,000	32	8	30	30	4	75,000/none	8/none
ecs.c7nex.16xlarge	64	128	80/none	8,000,000	32	15	50	50	4	150,000/none	16/none
ecs.c7nex.32xlarge	128	256	160/none	16,000,000	32	15	50	50	4	300,000/none	32/none

> **Note**
>
> Each ecs.c7nex.32xlarge instance must have at least two ENIs that are assigned different network card indexes before the instance can burst its network bandwidth to 160 Gbit/s. If all ENIs on the instance are assigned the same network card index, the instance can burst its network bandwidth only to 100 Gbit/s. For more information, see AttachNetworkInterface.

g7ne, network-enhanced general-purpose instance family

- **Introduction**: This instance family significantly improves the network throughput and packet forwarding rate per instance. A single instance can deliver a packet forwarding rate of up to 24,000,000 pps.
- **Supported scenarios**:
  - Network-intensive scenarios such as NFV or SD-WAN, mobile Internet, on-screen video comments, and telecom data forwarding
  - Small and medium-sized database systems, caches, and search clusters
  - Enterprise-level applications of various types and sizes
  - Big data analytics and machine learning
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses Intel® Xeon® Platinum 8369HB (Cooper Lake) or Intel® Xeon® Platinum 8369HC (Cooper Lake) processors that deliver a turbo frequency of 3.8 GHz and a clock speed of at least 3.3 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides high network performance based on large computing capacity.

The g7ne instance family includes the following instance types: ecs.g7ne.large, ecs.g7ne.xlarge, ecs.g7ne.2xlarge, ecs.g7ne.4xlarge, ecs.g7ne.8xlarge, ecs.g7ne.12xlarge, and ecs.g7ne.24xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.g7ne.large	2	8	1.5/10	900,000	450,000	2	3	10	10	10,000	0.75

ecs.g7ne.xlarge	4	16	3/10	1,000,000	900,000	4	4	15	15	20,000	1
ecs.g7ne.2xlarge	8	32	6/15	1,500,000	1,750,000	8	6	15	15	25,000	1.2
ecs.g7ne.4xlarge	16	64	12/25	3,000,000	3,500,000	16	8	30	30	40,000	2
ecs.g7ne.8xlarge	32	128	25/none	6,000,000	6,000,000	16	8	30	30	75,000	5
ecs.g7ne.12xlarge	48	192	40/none	12,000,000	8,000,000	32	8	30	30	100,000	8
ecs.g7ne.24xlarge	96	384	80/none	24,000,000	16,000,000	32	15	50	50	240,000	16

g5ne, network-enhanced general-purpose instance family

- **Introduction**: This instance family significantly improves the network throughput and packet forwarding rate per instance. A single instance can deliver a packet forwarding rate of up to 10,000,000 pps.
- **Supported scenarios**:
  - Data Plane Development Kit (DPDK) applications
  - Network-intensive scenarios such as NFV or SD-WAN, mobile Internet, on-screen video comments, and telecom data forwarding
  - Small and medium-sized database systems, caches, and search clusters
  - Enterprise-level applications of various types and sizes
  - Big data analytics and machine learning
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) or 8269CY (Cascade Lake) processors to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.

- Provides high network performance based on large computing capacity.

> **Note**
>
> To deploy DPDK applications, we recommend that you select instance types in the g5ne instance family.

The g5ne instance family includes the following instance types: ecs.g5ne.large, ecs.g5ne.xlarge, ecs.g5ne.2xlarge, ecs.g5ne.4xlarge, ecs.g5ne.8xlarge, ecs.g5ne.16xlarge, and ecs.g5ne.18xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.g5ne.large	2	8	1	400,000	450,000	2	3	10	10	10,000	1
ecs.g5ne.xlarge	4	16	2	750,000	900,000	4	4	15	15	15,000	1
ecs.g5ne.2xlarge	8	32	3.5	1,500,000	1,750,000	8	6	15	15	30,000	1
ecs.g5ne.4xlarge	16	64	7	3,000,000	3,500,000	16	8	30	30	60,000	2
ecs.g5ne.8xlarge	32	128	15	6,000,000	7,000,000	32	8	30	30	110,000	4
ecs.g5ne.16xlarge	64	256	30	12,000,000	14,000,000	32	8	30	30	130,000	8
ecs.g5ne.18xlarge	72	288	33	13,500,000	15,000,000	32	15	50	50	160,000	9

g7t, security-enhanced general-purpose instance family

- **Introduction**:

- This instance family supports up to 256 GiB of encrypted memory and confidential computing based on Intel® Software Guard Extensions (SGX) to protect the confidentiality and integrity of essential code and data from malware attacks.
- This instance family supports Virtual SGX (vSGX) and allows you to select instance types based on your business requirements.

  > **Important**
  >
  > If you use keys (such as SGX sealing keys) that are bound to hardware to encrypt the data of an instance within an Intel SGX enclave, the encrypted data cannot be decrypted after the host of the instance is changed. We recommend that you perform data redundancy and backup at the application layer to ensure application reliability.

- This instance family implements trusted boot based on Trusted Cryptography Module (TCM) or Trusted Platform Module (TPM) chips. During a trusted boot, all modules in the boot chain from the underlying server to the guest operating system are measured and verified.
- This instance family offloads a large number of virtualization features to dedicated hardware by using the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads.

- **Supported scenarios**:
  - Scenarios that involve sensitive information such as personal identity information, healthcare information, financial information, and intellectual property data
  - Scenarios where confidential data is shared among multiple parties
  - Blockchain scenarios
  - Confidential machine learning
  - Scenarios that require high security and enhanced trust, such as services for financial organizations, public service sectors, and enterprises
  - Enterprise-level applications of various types and sizes
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4. About 50% of memory is encrypted.
  - Uses the third-generation Intel® Xeon® Scalable (Ice Lake) processors that deliver a base frequency of 2.7 GHz and an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The g7t instance family includes the following instance types: ecs.g7t.large, ecs.g7t.xlarge, ecs.g7t.2xlarge, ecs.g7t.3xlarge, ecs.g7t.4xlarge, ecs.g7t.6xlarge, ecs.g7t.8xlarge, ecs.g7t.16xlarge, and ecs.g7t.32xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Encrypted memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Support for vTPM	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.g7t.large	2	8	4	2/burstable up to 10	900,000	Yes	Up to 250,000	2	3	6	6	20,000/burstable up to 110,000	1.5/burstable up to 6
ecs.g7t.xlarge	4	16	8	3/burstable up to 10	1,000,000	Yes	Up to 250,000	4	4	15	15	40,000/burstable up to 110,000	2/burstable up to 6
ecs.g7t.2xlarge	8	32	16	5/burstable up to 10	1,600,000	Yes	Up to 250,000	8	4	15	15	50,000/burstable up to 110,000	3/burstable up to 6
ecs.g7t.3xlarge	12	48	24	8/burstable up to 10	2,400,000	Yes	Up to 250,000	8	8	15	15	70,000/burstable up to 110,000	4/burstable up to 6
ecs.g7t.4xlarge	16	64	32	10/burstable up to 25	3,000,000	Yes	300,000	8	8	30	30	80,000/burstable up to 110,000	5/burstable up to 6
ecs.g7t.6xlarge	24	96	48	12/burstable up to 25	4,500,000	Yes	450,000	12	8	30	30	110,000/none	6/none

ecs.g7t.8xlarge	32	128	64	16/burstable up to 25	6,000,000	Yes	600,000	16	8	30	30	150,000/none	8/none
ecs.g7t.16xlarge	64	256	128	32/none	12,000,000	Yes	1,200,000	32	8	30	30	300,000/none	16/none
ecs.g7t.32xlarge	128	512	256	64/none	24,000,000	Yes	2,400,000	32	15	30	30	600,000/none	32/none

Note

- Intel Ice Lake supports only remote attestation based on Intel Software Guard Extensions Data Center Attestation Primitives (Intel SGX DCAP) and does not support remote attestation based on Intel Enhanced Privacy ID (EPID). You must adapt applications before you can use the remote attestation feature. For more information about remote attestation, see Strengthen Enclave Trust with Attestation.
- Intel SGX depends on host hardware. This instance family does not support hot migration.
- Operations, such as changing instance types and enabling the economical mode, may cause the host of an instance to change. For instances of this instance family, the host change may cause data decryption to fail. Proceed with caution.
- By default, failover is disabled. You can enable failover. For more information, see Modify instance maintenance attributes. Failover causes the host of an instance to change. For instances of this instance family, the host change may cause data decryption to fail. Proceed with caution.
- When you create a security-enhanced instance, you must select a dedicated image to use the security features. For more information, see Create a trusted instance.
- To use the ecs.g7t.32xlarge instance type, submit a ticket.

c7t, security-enhanced compute-optimized instance family

- **Introduction**:
  - This instance family supports up to 128 GiB of encrypted memory and confidential computing based on Intel® SGX to protect the confidentiality and integrity of essential code and data from malware attacks.
  - This instance family supports vSGX and allows you to select instance types based on your business requirements.

    Important

    If you use keys (such as SGX sealing keys) that are bound to hardware to encrypt the data of an instance within an Intel SGX enclave, the encrypted data cannot be decrypted after the host of the instance is changed. We recommend that you perform data redundancy and backup at the application layer to ensure application reliability.
  - This instance family implements trusted boot based on TCM or TPM chips. During a trusted boot, all modules in the boot chain from the underlying server to the guest operating system are measured and verified.
  - This instance family offloads a large number of virtualization features to dedicated hardware by using the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads.

- **Supported scenarios**:
  - Scenarios that involve sensitive information such as personal identity information, healthcare information, financial information, and intellectual property data
  - Scenarios where confidential data is shared among multiple parties
  - Blockchain scenarios
  - Confidential machine learning
  - Scenarios that require high security and enhanced trust, such as services for financial organizations, public service sectors, and enterprises
  - Enterprise-level applications of various types and sizes
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2. About 50% of memory is encrypted.
  - Uses the third-generation Intel® Xeon® Scalable (Ice Lake) processors that deliver a base frequency of 2.7 GHz and an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Provides high storage I/O performance based on large computing capacity.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The c7t instance family includes the following instance types: ecs.c7t.large, ecs.c7t.xlarge, ecs.c7t.2xlarge, ecs.c7t.3xlarge, ecs.c7t.4xlarge, ecs.c7t.6xlarge, ecs.c7t.8xlarge, ecs.c7t.16xlarge, and ecs.c7t.32xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Encrypted memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Support for vTPM	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.c7t.large	2	4	2	2/burstable up to 10	900,000	Yes	Up to 250,000	2	3	6	6	20,000/burstable up to 110,000	1.5/burstable up to 6

ecs.c7t.xlarge	4	8	4	3/burstable up to 10	1,000,000	Yes	Up to 250,000	4	4	15	15	40,000/burstable up to 110,000	2/burstable up to 6
ecs.c7t.2xlarge	8	16	8	5/burstable up to 10	1,600,000	Yes	Up to 250,000	8	4	15	15	50,000/burstable up to 110,000	3/burstable up to 6
ecs.c7t.3xlarge	12	24	12	8/burstable up to 10	2,400,000	Yes	Up to 250,000	8	8	15	15	70,000/burstable up to 110,000	4/burstable up to 6
ecs.c7t.4xlarge	16	32	16	10/burstable up to 25	3,000,000	Yes	300,000	8	8	30	30	80,000/burstable up to 110,000	5/burstable up to 6
ecs.c7t.6xlarge	24	48	24	12/burstable up to 25	4,500,000	Yes	450,000	12	8	30	30	110,000/none	6/none
ecs.c7t.8xlarge	32	64	32	16/burstable up to 25	6,000,000	Yes	600,000	16	8	30	30	150,000/none	8/none
ecs.c7t.16xlarge	64	128	64	32/none	12,000,000	Yes	1,200,000	32	8	30	30	300,000/none	16/none
ecs.c7t.32xlarge	128	256	128	64/none	24,000,000	Yes	2,400,000	32	15	30	30	600,000/none	32/none

> **Note**
>
> - Intel Ice Lake supports only remote attestation based on Intel Software Guard Extensions Data Center Attestation Primitives (Intel SGX DCAP) and does not support remote attestation based on Intel Enhanced Privacy ID (EPID). You must adapt applications before you can use the remote attestation feature. For more information about remote attestation, see Strengthen Enclave Trust with Attestation.
> - Intel SGX depends on host hardware. This instance family does not support hot migration.
> - Operations, such as changing instance types and enabling the economical mode, may cause the host of an instance to change. For instances of this instance family, the host change may cause data decryption to fail. Proceed with caution.
> - By default, failover is disabled. You can enable failover. For more information, see Modify instance maintenance attributes. Failover causes the host of an instance to change. For instances of this instance family, the host change may cause data decryption to fail. Proceed with caution.
> - When you create a security-enhanced instance, you must select a dedicated image to use the security features. For more information, see Create a trusted instance.
> - To use the ecs.c7t.32xlarge instance type, submit a ticket.

r7t, security-enhanced memory-optimized instance family

- **Introduction**:
  - This instance family supports up to 512 GiB of encrypted memory and confidential computing based on Intel® SGX to protect the confidentiality and integrity of essential code and data from malware attacks.
  - This instance family supports vSGX and allows you to select instance types that meet your business requirements.

    > **Important**
    >
    > If you use keys (such as SGX sealing keys) that are bound to hardware to encrypt the data of an instance within an Intel SGX enclave, the encrypted data cannot be decrypted after the host of the instance is changed. We recommend that you perform data redundancy and backup at the application layer to ensure application reliability.

  - This instance family implements trusted boot based on TCM or TPM chips. During a trusted boot, all modules in the boot chain from the underlying server to the guest operating system are measured and verified.
  - This instance family offloads a large number of virtualization features to dedicated hardware by using the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads.
- **Supported scenarios**:
  - Encrypted computing applications for databases
  - Scenarios that involve sensitive information such as personal identity information, healthcare information, financial information, and intellectual property data
  - Scenarios where confidential data is shared among multiple parties
  - Blockchain scenarios
  - Confidential machine learning
  - Scenarios that require high security and enhanced trust, such as services for financial organizations, public service sectors, and enterprises
  - Enterprise-level applications of various types and sizes
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8. About 50% of memory is encrypted.

- Uses the third-generation Intel® Xeon® Scalable (Ice Lake) processors that deliver a base frequency of 2.7 GHz and an all-core turbo frequency of 3.5 GHz to provide consistent computing performance.
- Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Provides high storage I/O performance based on large computing capacity.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The r7t instance family includes the following instance types: ecs.r7t.large, ecs.r7t.xlarge, ecs.r7t.2xlarge, ecs.r7t.3xlarge, ecs.r7t.4xlarge, ecs.r7t.6xlarge, ecs.r7t.8xlarge, ecs.r7t.16xlarge, and ecs.r7t.32xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Encrypted memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Support for vTPM	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.r7t.large	2	16	8	2/burstable up to 10	900,000	Yes	Up to 250,000	2	3	6	6	20,000/burstable up to 110,000	1.5/burstable up to 6
ecs.r7t.xlarge	4	32	16	3/burstable up to 10	1,000,000	Yes	Up to 250,000	4	4	15	15	40,000/burstable up to 110,000	2/burstable up to 6
ecs.r7t.2xlarge	8	64	32	5/burstable up to 10	1,600,000	Yes	Up to 250,000	8	4	15	15	50,000/burstable up to 110,000	3/burstable up to 6

ecs.r7t.3xlarge	12	96	48	8/burstable up to 10	2,400,000	Yes	Up to 250,000	8	8	15	15	70,000/burstable up to 110,000	4/burstable up to 6
ecs.r7t.4xlarge	16	128	64	10/burstable up to 25	3,000,000	Yes	300,000	8	8	30	30	80,000/burstable up to 110,000	5/burstable up to 6
ecs.r7t.6xlarge	24	192	96	12/burstable up to 25	4,500,000	Yes	450,000	12	8	30	30	110,000/none	6/none
ecs.r7t.8xlarge	32	256	128	16/burstable up to 25	6,000,000	Yes	600,000	16	8	30	30	150,000/none	8/none
ecs.r7t.16xlarge	64	512	256	32/none	12,000,000	Yes	1,200,000	32	8	30	30	300,000/none	16/none
ecs.r7t.32xlarge	128	1024	512	64/none	24,000,000	Yes	2,400,000	32	15	30	30	600,000/none	32/none

Note

- Intel Ice Lake supports only remote attestation based on Intel Software Guard Extensions Data Center Attestation Primitives (Intel SGX DCAP) and does not support remote attestation based on Intel Enhanced Privacy ID (EPID). You must adapt applications before you can use the remote attestation feature. For more information about remote attestation, see Strengthen Enclave Trust with Attestation.
- Intel SGX depends on host hardware. This instance family does not support hot migration.
- Operations, such as changing instance types and enabling the economical mode, may cause the host of an instance to change. For instances of this instance family, the host change may cause data decryption to fail. Proceed with caution.
- By default, failover is disabled. You can enable failover. For more information, see Modify instance maintenance attributes. Failover causes the host of an instance to change. For instances of this instance family, the host change may cause data decryption to fail. Proceed with caution.
- When you create a security-enhanced instance, you must select a dedicated image to use the security features. For more information, see Create a trusted instance.
- To use the ecs.r7t.32xlarge instance type, submit a ticket.

g6t, security-enhanced general-purpose instance family

Features:

- **Introduction**:
  - This instance family implements trusted boot based on TCM or TPM chips. During a trusted boot, all modules in the boot chain from the underlying server to the guest operating system are measured and verified.
  - This instance family supports the vTPM feature and delivers ~~a full set of~~ trusted capabilities at the IaaS layer based on integrity monitoring.
  - This instance family offloads a large number of virtualization features to dedicated hardware by using the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**:
  - Scenarios that require high security and enhanced trust, such as services for financial organizations, public service sectors, and enterprises
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Enterprise-level applications of various types and sizes
  - Websites and application servers
  - Game servers
  - Small and medium-sized database systems, caches, and search clusters
  - Data analytics and computing
  - Computing clusters and memory-intensive data processing
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver a turbo frequency of 3.2 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The g6t instance family includes the following instance types: ecs.g6t.large, ecs.g6t.xlarge, ecs.g6t.2xlarge, ecs.g6t.4xlarge, ecs.g6t.8xlarge, ecs.g6t.13xlarge, and ecs.g6t.26xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Support for vTPM	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.g6t.large	2	8	1.2/burstable up to 10	900,000	Yes	Up to 250,000	2	3	6	1	20,000	1
ecs.g6t.xlarge	4	16	2/burstable up to 10	1,000,000	Yes	Up to 250,000	4	4	15	1	40,000	1.5
ecs.g6t.2xlarge	8	32	3/burstable up to 10	1,600,000	Yes	Up to 250,000	8	4	15	1	50,000	2
ecs.g6t.4xlarge	16	64	6/burstable up to 10	3,000,000	Yes	300,000	8	8	30	1	80,000	3
ecs.g6t.8xlarge	32	128	10/none	6,000,000	Yes	600,000	16	8	30	1	150,000	5
ecs.g6t.13xlarge	52	192	16/none	9,000,000	Yes	900,000	32	7	30	1	240,000	8
ecs.g6t.26xlarge	104	384	32/none	24,000,000	Yes	1,800,000	32	15	30	1	480,000	16

Note

The results for network capabilities are the maximum values obtained from single-item tests. For example, when network bandwidth is tested, no stress tests are performed on the packet forwarding rate or other network metrics.

c6t, security-enhanced compute-optimized instance family

- **Introduction**:
  - This instance family implements trusted boots based on TPM chips. During a trusted boot, all modules in the boot chain from the underlying hardware to the guest operating system are measured and verified.
  - This instance family supports integrity monitoring and provides a full set of trusted capabilities at the IaaS layer.
  - This instance family offloads a large number of virtualization features to dedicated hardware by using the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance and reduce virtualization overheads. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**:
  - Scenarios that require high security and enhanced trust, such as services for financial organizations, public service sectors, and enterprises
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Web frontend servers
  - Frontend servers of massively multiplayer online (MMO) games
  - Data analytics, batch processing, and video encoding
  - High-performance scientific and engineering applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver a turbo frequency of 3.2 GHz to provide consistent computing performance.
  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides high network performance based on large computing capacity.

The c6t instance family includes the following instance types: ecs.c6t.large, ecs.c6t.xlarge, ecs.c6t.2xlarge, ecs.c6t.4xlarge, ecs.c6t.8xlarge, ecs.c6t.13xlarge, and ecs.c6t.26xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Support for vTPM	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.c6t.large	2	4	1.2/burstable up to 10	900,000	Yes	Up to 250,000	2	3	6	1	20,000	1
ecs.c6t.xlarge	4	8	2/burstable up to 10	1,000,000	Yes	Up to 250,000	4	4	15	1	40,000	1.5
ecs.c6t.2xlarge	8	16	3/burstable up to 10	1,600,000	Yes	Up to 250,000	8	4	15	1	50,000	2
ecs.c6t.4xlarge	16	32	6/burstable up to 10	3,000,000	Yes	300,000	8	8	30	1	80,000	3
ecs.c6t.8xlarge	32	64	10/none	6,000,000	Yes	600,000	16	8	30	1	150,000	5
ecs.c6t.13xlarge	52	96	16/none	9,000,000	Yes	900,000	32	7	30	1	240,000	8
ecs.c6t.26xlarge	104	192	32/none	24,000,000	Yes	1,800,000	32	15	30	1	480,000	16

> **Note**
>
> The results for network capabilities are the maximum values obtained from single-item tests. For example, when network bandwidth is tested, no stress tests are performed on the packet forwarding rate or other network metrics.

re6p, persistent memory-optimized instance family

For answers to commonly asked questions about persistent memory-optimized instances, see Instance FAQ.

Features:

- **Introduction**:
  - This instance family uses Intel® Optane™ persistent memory.

    > **Important**
    >
    > The reliability of data stored in persistent memory varies based on the reliability of persistent memory devices and the physical servers to which these devices are attached. Risks of single points of failure exist. To ensure the reliability of application data, we recommend that you implement data redundancy at the application layer and use cloud disks for long-term data storage.

  - This instance family allows persistent memory to be used as memory or as local SSDs on instances of some instance types.

    > **Note**
    >
    > For more information, see Configure the usage mode of persistent memory.

  - This instance family provides the ecs.re6p-redis.<nx>large instance types for Redis applications.

    > **Note**
    >
    > ecs.re6p-redis.<nx>large instance types are exclusively provided for Redis applications. Persistent memory on instances of these instance types is used as memory by default and cannot be re-configured as local SSDs. For information about how to deploy a Redis application, see Deploy Redis on persistent memory-optimized instances.

- **Supported scenarios**:
  - Redis and other NoSQL databases such as Cassandra and MongoDB
  - Structured databases such as MySQL
  - I/O-intensive applications such as e-commerce, online games, and media applications
  - Search scenarios that use solutions such as Elasticsearch
  - Live video streaming, instant messaging, and room-based online games that require persistent connections
  - High-performance relational databases and OLTP systems
- **Compute**:
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver a turbo frequency of 3.2 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:

- Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.

The re6p instance family includes the following instance types: ecs.re6p.large, ecs.re6p.xlarge, ecs.re6p.2xlarge, ecs.re6p.13xlarge, ecs.re6p.26xlarge, ecs.re6p-redis.large, ecs.re6p-redis.xlarge, ecs.re6p-redis.2xlarge, and ecs.re6p-redis.13xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Persistent memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.re6p.large	2	8	31.5	1/3	300,000	Up to 250,000	2	2	6	1	10,000	1
ecs.re6p.xlarge	4	16	63	1.5/5	500,000	Up to 250,000	4	3	10	1	20,000	1.5
ecs.re6p.2xlarge	8	32	126	2.5/10	800,000	Up to 250,000	8	4	20	1	25,000	2
ecs.re6p.13xlarge	52	192	756	12.5/none	3,000,000	900,000	32	7	20	1	100,000	8
ecs.re6p.26xlarge	104	384	1512	25/none	6,000,000	1,800,000	32	15	20	1	200,000	16,0
ecs.re6p-redis.large	2	8	31.5	1/3	300,000	Up to 250,000	2	2	6	1	10,000	1
ecs.re6p-redis.xlarge	4	16	63	1.5/5	500,000	Up to 250,000	4	3	10	1	20,000	1.5
ecs.re6p-redis.2xlarge	8	32	126	2.5/10	800,000	Up to 250,000	8	4	20	1	25,000	2

ecs.re6p-redis.13xlarge	52	192	756	12.5/none	3,000,000	900,000	32	7	20	1	100,000	8

re6, high-memory instance family

Features:

- **Introduction**: This instance family is optimized for high-performance databases, in-memory databases, and enterprise-level memory-intensive applications.
- **Supported scenarios**:
  - High-performance databases and in-memory databases such as SAP HANA
  - Memory-intensive applications
  - Big data processing engines such as Apache Spark and Presto
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:15 and up to 3 TiB of memory.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver a turbo frequency of 3.2 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.

The re6 instance family includes the following instance types: ecs.re6.4xlarge, ecs.re6.8xlarge, ecs.re6.13xlarge, ecs.re6.16xlarge, ecs.re6.26xlarge, ecs.re6.32xlarge, and ecs.re6.52xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.re6.4xlarge	16	256	5	900,000	8	7	20	1	25,000	2
ecs.re6.8xlarge	32	512	10	1,800,000	16	7	20	1	50,000	4
ecs.re6.13xlarge	52	768	10	1,800,000	16	7	20	1	50,000	4

ecs.re6.16xlarge	64	1024	16	3,000,000	32	7	20	1	100,000	8
ecs.re6.26xlarge	104	1536	16	3,000,000	32	7	20	1	100,000	8
ecs.re6.32xlarge	128	2048	32	6,000,000	32	15	20	1	200,000	16
ecs.re6.52xlarge	208	3072	32	6,000,000	32	15	20	1	200,000	16

Note

To use the ecs.re6.32xlarge instance type, submit a ticket.

re4, high-memory instance family

- **Introduction**:
  - This instance family is optimized for high-performance databases, in-memory databases, and enterprise-level memory-intensive applications.
  - The ecs.re4.20xlarge and ecs.re4.40xlarge instance types are SAP HANA-certified.
- **Supported scenarios**:
  - High-performance databases and in-memory databases such as SAP HANA
  - Memory-intensive applications
  - Big data processing engines such as Apache Spark and Presto
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:12 and up to 1,920 GiB of memory.
  - Uses 2.2 GHz Intel® Xeon® E7 8880 v4 (Broadwell) processors that deliver a turbo frequency of up to 2.4 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.

The re4 instance family includes the following instance types: ecs.re4.10xlarge, ecs.re4.20xlarge, and ecs.re4.40xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI

ecs.re4.10xlarge	40	480	8	1,000,000	8	4	10	1
ecs.re4.20xlarge	80	960	15	2,000,000	16	8	20	1
ecs.re4.40xlarge	160	1920	30	4,500,000	16	8	20	1

re4e, high-memory instance family

To use the re4e instance family, submit a ticket.

- **Introduction**: This instance family is optimized for high-performance databases, in-memory databases, and enterprise-level memory-intensive applications.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:24 and up to 3,840 GiB of memory.
  - Uses 2.2 GHz Intel® Xeon® E7 8880 v4 (Broadwell) processors that deliver a turbo frequency of up to 2.4 GHz to provide consistent computing performance.
- **Supported scenarios**:
  - High-performance databases and in-memory databases such as SAP HANA
  - Memory-intensive applications
  - Big data processing engines such as Apache Spark and Presto
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.

Instance types

Instance type	vCPU	Memory (GiB)	Network bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	Private IPv6 addresses per ENI
ecs.re4e.40xlarge	160	3840	30	4,500,000	16	15	20	1

x86-based entry-level computing instance families

e, economy instance family

Features:

- Compute:
  - Offers multiple CPU-to-memory ratios such as 1:1, 1:2, and 1:4.

- Uses Intel® Xeon® Platinum Scalable processors.

> **Note**
>
> Instances of the e instance family use a CPU-unbound scheduling scheme, in which each vCPU is randomly allocated to an idle CPU hyperthread. Compared with enterprise-level instances, e instances share resources and cost less.

- Storage:
  - Is an instance family in which all instances are I/O optimized.
  - Supports enhanced SSDs (ESSDs), ESSD Entry disks, and ESSD AutoPL disks.

> **Note**
>
> Due to the limits of economy instance types, ESSDs at performance levels 1, 2, and 3 (PL1, PL2, and PL3 ESSDs) cannot deliver their maximum performance on e instances. We recommend that you select ESSD Entry disks or PL0 ESSDs for the instances.

- Network:
  - Supports IPv4 and IPv6.
  - Supports only virtual private clouds (VPCs).
  - Provides high network performance based on large computing capacity.
- Supported scenarios:
  - Small and medium-sized websites
  - Development and testing
  - Lightweight applications

Instance types

Instance type	vCPUs	Memory size (GiB)	Baseline/burst bandwidth (Gbit/s)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.e-c4m1.large	2	0.5	0.2/burstable up to 2	1	2	2	1	8,000/none	0.4/none
ecs.e-c2m1.large	2	1	0.2/burstable up to 2	1	2	2	1	8,000/none	0.4/none
ecs.e-c1m1.large	2	2.0	0.2/burstable up to 2	1	2	2	1	8,000/none	0.4/none
ecs.e-c1m2.large	2	4.0	0.2/burstable up to 2	1	2	2	1	8,000/none	0.4/none
ecs.e-c1m4.large	2	8.0	0.4/burstable up to 2	1	2	2	1	16,000/none	0.8/none

ecs.e-c1m2.xlarge	4	8.0	0.4/burstable up to 3	1	2	6	1	16,000/none	0.8/none
ecs.e-c1m4.xlarge	4	16.0	0.8/burstable up to 4	1	2	6	1	16,000/none	0.8/none
ecs.e-c1m2.2xlarge	8	16.0	0.8/burstable up to 6	1	2	6	1	16,000/none	0.8/none
ecs.e-c1m4.2xlarge	8	32.0	1.2/burstable up to 6	1	2	6	1	16,000/none	0.8/none

Note

- You can go to the Instance Types Available for Each Region page to view the instance types available in each region.
- For more information about these specifications, see the "Instance type specifications" section in Overview of instance families. Packet forwarding rates vary significantly based on business scenarios. We recommend that you perform business stress tests on instances to choose appropriate instance types.
- The following limits apply to the ecs.e-c4m1.large, ecs.e-c2m1.large, ecs.e-c1m1.large, ecs.e-c1m2.large, and ecs.e-c1m4.large instance types:
  - Secondary elastic network interfaces (ENIs) cannot be bound to ecs.e-c1m1.large, ecs.e-c1m2.large, or ecs.e-c1m4.large instances during instance creation and can be bound after the instances are created.
  - You can bind secondary ENIs to or unbind secondary ENIs from ecs.e-c1m1.large, ecs.e-c1m2.large, and ecs.e-c1m4.large instances only when the instances are in the Stopped state.
- The ecs.e-c4m1.large and ecs.e-c2m1.large1.large instance types are available for purchase only in the following regions: China (Hong Kong), Singapore, Malaysia (Kuala Lumpur), Indonesia (Jakarta), Philippines (Manila), Thailand (Bangkok), Japan (Tokyo), South Korea (Seoul), UK (London), Germany (Frankfurt), US (Virginia), and US (Silicon Valley).

t6, burstable instance family

Features:

- Provides a CPU performance baseline and the ability to burst above the baseline, which are governed by accrued CPU credits.
- More cost-effective compared with the t5 burstable instance family.
- Compute:
  - Uses 2.5 GHz Intel® Xeon® Cascade Lake processors that deliver a turbo frequency of 3.2 GHz.
  - Uses DDR4 memory.
- Storage:
  - Is an I/O optimized instance.

- Supports Enterprise SSDs (ESSDs), ESSD AutoPL disks, standard SSDs, and ultra disks.

  > **Important**
  >
  > ESSDs at performance level (PL) 2 and 3 cannot provide maximum performance due to the specification limits of burstable instances. We recommend that you use enterprise-level instances or ESSDs of lower performance levels.

- Network:
  - Supports IPv4 and IPv6.
  - Supports only virtual private clouds (VPCs).
- Supported scenarios:
  - Web application servers
  - Lightweight applications and microservices
  - Development and testing environments

Instance types

Instance type	vCPU	Memory (GiB)	Average baseline CPU performance	CPU credits per hour	Max CPU credit balance	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.t6-c4m1.large	2	0.5	5%	6	144	0.08/burstable up to 0.4	40,000	1	2	2	1
ecs.t6-c2m1.large	2	1.0	10%	12	288	0.08/burstable up to 0.6	60,000	1	2	2	1
ecs.t6-c1m1.large	2	2.0	20%	24	576	0.08/burstable up to 1	100,000	1	2	2	1
ecs.t6-c1m2.large	2	4.0	20%	24	576	0.08/burstable up to 1	100,000	1	2	2	1
ecs.t6-c1m4.large	2	8.0	30%	36	864	0.08/burstable up to 1	100,000	1	2	2	1
ecs.t6-c1m4.xlarge	4	16.0	40%	96	2304	0.16/burstable up to 2	200,000	1	2	6	1

ecs.t6-c1m4.2xlarge	8	32.0	40%	192	4608	0.32/burstable up to 4	400,000	1	2	6	1

Note

- Secondary elastic network interfaces (ENIs) cannot be bound to instances of this instance family when the instances are being created and can be bound to the instances after the instances are created. When you bind secondary ENIs to or unbind secondary ENIs from instances of the following instance types, make sure that the instances are in the Stopped state: ecs.t6-c1m1.large, ecs.t6-c1m2.large, ecs.t6-c1m4.large, ecs.t6-c2m1.large, and ecs.t6-c4m1.large.
- You can go to the Instance Types Available for Each Region page to view the instance types available in each region.
- For information about instance type metrics, see Instance type metrics.

t5, burstable instance family

Features:

- Provides a CPU performance baseline and the ability to burst above the baseline, which are governed by accrued CPU credits.
- Balances compute, memory, and network resources.
- Compute:
  - Offers multiple CPU-to-memory ratios.
  - Uses 2.5 GHz Intel® Xeon® processors.
  - Uses DDR4 memory.
- Storage: supports only ultra disks and standard SSDs.
- Network:
  - Supports IPv4 and IPv6.
  - Supports only VPCs.
- Supported scenarios:
  - Web application servers
  - Lightweight applications and microservices
  - Development and testing environments

Instance types

Instance type	vCPU	Memory (GiB)	Average baseline CPU performance	CPU credits per hour	Max CPU credit balance	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.t5-lc2m1.nano	1	0.5	20%	12	288	0.1	40,000	1	2	2	1

ecs.t5-lc1m1.small	1	1.0	20%	12	288	0.2	60,000	1	2	2	1
ecs.t5-lc1m2.small	1	2.0	20%	12	288	0.2	60,000	1	2	2	1
ecs.t5-lc1m2.large	2	4.0	20%	24	576	0.4	100,000	1	2	2	1
ecs.t5-lc1m4.large	2	8.0	20%	24	576	0.4	100,000	1	2	2	1
ecs.t5-c1m1.large	2	2.0	25%	30	720	0.5	100,000	1	2	2	1
ecs.t5-c1m2.large	2	4.0	25%	30	720	0.5	100,000	1	2	2	1
ecs.t5-c1m4.large	2	8.0	25%	30	720	0.5	100,000	1	2	2	1
ecs.t5-c1m1.xlarge	4	4.0	25%	60	1440	0.8	200,000	1	2	6	1
ecs.t5-c1m2.xlarge	4	8.0	25%	60	1440	0.8	200,000	1	2	6	1
ecs.t5-c1m4.xlarge	4	16.0	25%	60	1440	0.8	200,000	1	2	6	1
ecs.t5-c1m1.2xlarge	8	8.0	25%	120	2880	1.2	400,000	1	2	6	1
ecs.t5-c1m2.2xlarge	8	16.0	25%	120	2880	1.2	400,000	1	2	6	1
ecs.t5-c1m4.2xlarge	8	32.0	25%	120	2880	1.2	400,000	1	2	6	1

ecs.t5-c1m1.4xlarge	16	16.0	25%	240	5760	1.2	600,000	1	2	6	1
ecs.t5-c1m2.4xlarge	16	32.0	25%	240	5760	1.2	600,000	1	2	6	1

Note

- Secondary ENIs cannot be bound to instances of this instance family when the instances are being created and can be bound to the instances after the instances are created. When you bind secondary ENIs to or unbind secondary ENIs from instances of the following instance types, make sure that the instances are in the Stopped state: ecs.t5-lc2m1.nano, ecs.t5-c1m1.large, ecs.t5-c1m2.large, ecs.t5-c1m4.large, ecs.t5-lc1m1.small, ecs.t5-lc1m2.large, ecs.t5-lc1m2.small, and ecs.t5-lc1m4.large.
- You can go to the Instance Types Available for Each Region page to view the instance types available in each region.
- For information about instance type metrics, see Instance type metrics.

v5, CPU overprovisioned instance family

Features

- You can create v5 instances only on dedicated hosts.
- Compute:
  - Supports multiple CPU-to-memory ratios such as 1:1, 1:2, 1:4, and 1:8.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors.
- Storage:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, standard SSDs, and ultra disks.
- Network:
  - Supports IPv6.
- Suits the following scenarios:
  - Migration from offline virtualization environments to Alibaba Cloud
  - Services that generate low, medium, or burstable CPU loads

Instance types

Instance type	vCPUs	Memory (GiB)	Bandwidth (Gbit/s)	Packet forwarding rate (Kpps)	NIC queues	ENIs	Private IP addresses per ENI
ecs.v5-c1m1.large	2	2.0	2.0	300	2	2	2
ecs.v5-c1m1.xlarge	4	4.0	2.0	300	2	2	6
ecs.v5-c1m1.2xlarge	8	8.0	3.0	400	2	3	6
ecs.v5-c1m1.3xlarge	12	12.0	3.0	400	4	3	6

ecs.v5-c1m1.4xlarge	16	16.0	4.0	500	4	4	6
ecs.v5-c1m1.8xlarge	32	32.0	4.0	500	8	4	6
ecs.v5-c1m2.large	2	4.0	2.0	300	2	2	2
ecs.v5-c1m2.xlarge	4	8.0	2.0	300	2	2	6
ecs.v5-c1m2.2xlarge	8	16.0	3.0	400	2	3	6
ecs.v5-c1m2.3xlarge	12	24.0	3.0	400	4	3	6
ecs.v5-c1m2.4xlarge	16	32.0	4.0	500	4	4	6
ecs.v5-c1m2.8xlarge	32	64.0	4.0	500	8	4	6
ecs.v5-c1m4.large	2	8.0	2.0	300	2	2	2
ecs.v5-c1m4.xlarge	4	16.0	2.0	300	2	2	6
ecs.v5-c1m4.2xlarge	8	32.0	3.0	400	2	3	6
ecs.v5-c1m4.3xlarge	12	48.0	3.0	400	4	3	6
ecs.v5-c1m4.4xlarge	16	64.0	4.0	500	4	4	6
ecs.v5-c1m4.8xlarge	32	128.0	4.0	500	8	4	6
ecs.v5-c1m8.large	2	16.0	2.0	300	2	2	2
ecs.v5-c1m8.xlarge	4	32.0	2.0	300	2	2	6
ecs.v5-c1m8.2xlarge	8	64.0	3.0	400	2	3	6

ecs.v5-c1m8.3xlarge	12	96.0	3.0	400	4	3	6
ecs.v5-c1m8.4xlarge	16	128.0	4.0	500	4	4	6
ecs.v5-c1m8.8xlarge	32	256.0	4.0	500	8	4	6

Note For more information about these specifications, see Overview of instance families.

xn4, n4, mn4, and e4, previous-generation shared instance families

Features:

- Offer multiple CPU-to-memory ratios.
- Use 2.5 GHz Intel® Xeon® processors.
- Use DDR4 memory.
- Are instance families in which all instances are I/O optimized.
- Support only IPv4.

Instance family	Description	vCPU-to-memory ratio	Scenario
xn4	Shared compact instance family	1:1	• Web frontend applications • Lightweight applications and microservices • Development and testing environments
n4	Shared compute instance family	1:2	• Websites and web applications • Development environments, servers, code repositories, microservices, and testing and staging environments • Lightweight enterprise-level applications
mn4	Shared general-purpose instance family	1:4	• Websites and web applications • Lightweight databases and caches • Integrated applications and lightweight enterprise-level services

e4	Shared memory instance family	1:8	• Applications that require a large amount of memory • Lightweight databases and caches

xn4 instance types

Instance type	vCPUs	Memory size (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Network interface controller (NIC) queues	ENIs	Private IPv4 addresses per ENI
ecs.xn4.small	1	1.0	0.5	5	1	2	2

Note

- Secondary ENIs cannot be bound to instances of this instance family during instance creation and can be bound after the instances are created. You can bind secondary ENIs to or unbind secondary ENIs from an ecs.xn4.small instance only when the instance is in the Stopped state.
- You can go to the Instance Types Available for Each Region page to view the instance types available in each region.
- For more information about these specifications, see the "Instance type specifications" section in Overview of instance families. Packet forwarding rates vary significantly based on business scenarios. We recommend that you perform business stress tests on instances to choose appropriate instance types.

n4 instance types

Instance type	vCPUs	Memory size (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI
ecs.n4.small	1	2.0	0.5	5	1	2	2
ecs.n4.large	2	4.0	0.5	10	1	2	2
ecs.n4.xlarge	4	8.0	0.8	15	1	2	6
ecs.n4.2xlarge	8	16.0	1.2	30	1	2	6
ecs.n4.4xlarge	16	32.0	2.5	40	1	2	6
ecs.n4.8xlarge	32	64.0	5.0	50	1	2	6

Note

- Secondary ENIs cannot be bound to instances of this instance family during instance creation and can be bound after the instances are created. You can bind secondary ENIs to or unbind secondary ENIs from instances of specific instance types, including ecs.n4.small and ecs.n4.large, only when the instances are in the Stopped state.
- You can go to the Instance Types Available for Each Region page to view the instance types available in each region.
- For more information about these specifications, see the "Instance type specifications" section in Overview of instance families. Packet forwarding rates vary significantly based on business scenarios. We recommend that you perform business stress tests on instances to choose appropriate instance types.

mn4 instance types

Instance type	vCPUs	Memory size (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI
ecs.mn4.small	1	4.0	0.5	5	1	2	2
ecs.mn4.large	2	8.0	0.5	10	1	2	2
ecs.mn4.xlarge	4	16.0	0.8	15	1	2	6
ecs.mn4.2xlarge	8	32.0	1.2	30	1	2	6
ecs.mn4.4xlarge	16	64.0	2.5	40	1	8	6
ecs.mn4.8xlarge	32	128.0	5	50	2	8	6

Note

- Secondary ENIs cannot be bound to instances of this instance family during instance creation and can be bound after the instances are created. You can bind secondary ENIs to or unbind secondary ENIs from instances of specific instance types, including ecs.mn4.small and ecs.mn4.large, only when the instances are in the Stopped state.
- You can go to the Instance Types Available for Each Region page to view the instance types available in each region.
- For more information about these specifications, see the "Instance type specifications" section in Overview of instance families. Packet forwarding rates vary significantly based on business scenarios. We recommend that you perform business stress tests on instances to choose appropriate instance types.

e4 instance types

Instance type	vCPUs	Memory size (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI
ecs.e4.small	1	8.0	0.5	5	1	2	2
ecs.e4.large	2	16.0	0.5	10	1	2	2
ecs.e4.xlarge	4	32.0	0.8	15	1	2	6
ecs.e4.2xlarge	8	64.0	1.2	30	1	3	6
ecs.e4.4xlarge	16	128.0	2.5	40	1	8	6

Note

- Secondary ENIs cannot be bound to instances of this instance family during instance creation and can be bound after the instances are created. You can bind secondary ENIs to or unbind secondary ENIs from instances of specific instance types, including ecs.e4.small and ecs.e4.large, only when the instances are in the Stopped state.
- You can go to the Instance Types Available for Each Region page to view the instance types available in each region.
- For more information about these specifications, see the "Instance type specifications" section in Overview of instance families. Packet forwarding rates vary significantly based on business scenarios. We recommend that you perform business stress tests on instances to choose appropriate instance types.

Arm-based enterprise-level computing instance families

g8y, general-purpose instance family

- **Introduction**: This instance family uses in-house Arm-based YiTian 710 processors and the fourth-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**: containers, microservices, websites, application servers, video encoding and decoding, HPC, and CPU-based machine learning.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.75 GHz Yitian 710 processors to provide consistent computing performance.
- **Storage**:

- Is an instance family in which all instances are I/O optimized.
- Supports the Non-Volatile Memory Express (NVMe) protocol. For more information, see NVMe protocol.
- Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- Provides high network and storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports elastic RDMA interfaces (ERIs). For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.

g8y instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Maximum attached data disks	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.g8y.small	1	4	1/10	500,000	Up to 250,000	1	2	3	3	5	10,000/burstable up to 110,000	1/burstable up to 10
ecs.g8y.large	2	8	2/10	900,000	Up to 250,000	2	3	6	6	8	20,000/burstable up to 110,000	1.5/burstable up to 10
ecs.g8y.xlarge	4	16	3/10	1,000,000	Up to 250,000	4	4	15	15	8	40,000/burstable up to 110,000	2/burstable up to 10
ecs.g8y.2xlarge	8	32	5/10	1,600,000	Up to 250,000	8	4	15	15	16	50,000/burstable up to 110,000	3/burstable up to 10

ecs.g8y.4xlarge	16	64	10/25	3,000,000	400,000	25	8	30	30	16	80,000/burstable up to 110,000	5/burstable up to 10
ecs.g8y.8xlarge	32	128	16/25	5,000,000	750,000	32	8	30	30	16	125,000	10
ecs.g8y.16xlarge	64	256	32/none	10,000,000	1,500,000	32	8	30	30	32	250,000	16
ecs.g8y.32xlarge	128	512	64/none	20,000,000	3,000,000	32	15	30	30	32	500,000	32

Note

If you want to use the ecs.g8y.32xlarge instance type, submit a ticket.

c8y, compute-optimized instance family

- **Introduction**: This instance family uses in-house Arm-based YiTian 710 processors and the fourth-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**: containers, microservices, websites, application servers, video encoding and decoding, HPC, and CPU-based machine learning.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.75 GHz YiTian 710 processors to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports the NVMe protocol. For more information, see NVMe protocol.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Offers burstable disk IOPS and burstable disk bandwidth for low-specification instances and provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports elastic RDMA interfaces (ERIs). For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides ultra-high packet forwarding rates.

- Provides burstable network bandwidth for low-specification instances.
- Provides high network performance based on large computing capacity.

The c8y instance family includes the following instance types: ecs.c8y.small, ecs.c8y.large, ecs.c8y.xlarge, ecs.c8y.2xlarge, ecs.c8y.4xlarge, ecs.c8y.8xlarge, ecs.c8y.16xlarge, and ecs.c8y.32xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	ERIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Maximum attached data disks	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.c8y.small	1	2	1/10	500,000	Up to 250,000	1	2	0	3	3	5	10,000/burstable up to 110,000	1/burstable up to 6
ecs.c8y.large	2	4	2/10	900,000	Up to 250,000	2	3	1	6	6	8	20,000/burstable up to 110,000	1.5/burstable up to 6
ecs.c8y.xlarge	4	8	3/10	1,000,000	Up to 250,000	4	4	1	15	15	8	40,000/burstable up to 110,000	2/burstable up to 6
ecs.c8y.2xlarge	8	16	5/10	1,600,000	Up to 250,000	8	4	1	15	15	16	50,000/burstable up to 110,000	3/burstable up to 6
ecs.c8y.4xlarge	16	32	10/25	3,000,000	400,000	25	8	1	30	30	16	80,000/burstable up to 110,000	5/burstable up to 6

ecs.c8y.8xlarge	32	64	16/none	5,000,000	750,000	32	8	1	30	30	16	125,000	8
ecs.c8y.16xlarge	64	128	32/none	10,000,000	1,500,000	32	8	1	30	30	32	250,000	16
ecs.c8y.32xlarge	128	256	64/none	20,000,000	3,000,000	32	15	1	30	30	32	500,000	32

Note

If you want to use the ecs.c8y.32xlarge instance type, submit a ticket.

r8y, memory-optimized instance family

- **Introduction**: This instance family uses in-house Arm-based YiTian 710 processors and the fourth-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**: scenarios such as containers, microservices, websites and application servers, video encoding and decoding, high-performance computing, and CPU-based machine learning.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses 2.75 GHz YiTian 710 processors to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports the NVMe protocol. For more information, see NVMe protocol.
  - Supports ESSDs and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
  - Provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports ERIs. For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides ultra-high packet forwarding rates.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.

The r8y instance family includes the following instance types: ecs.r8y.small, ecs.r8y.large, ecs.r8y.xlarge, ecs.r8y.2xlarge, ecs.r8y.4xlarge, ecs.r8y.8xlarge, ecs.r8y.16xlarge, and ecs.r8y.32xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see the Instance type metrics section of the "Classification and naming of instance types" topic.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	ERIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Maximum attached data disks	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.r8y.small	1	8	1/10	500,000	Up to 250,000	1	2	0	3	3	5	10,000/burstable up to 110,000	1/burstable up to 6
ecs.r8y.large	2	16	2/10	900,000	Up to 250,000	2	3	1	6	6	8	20,000/burstable up to 110,000	1.5/burstable up to 6
ecs.r8y.xlarge	4	32	3/10	1,000,000	Up to 250,000	4	4	1	15	15	8	40,000/burstable up to 110,000	2/burstable up to 6
ecs.r8y.2xlarge	8	64	5/10	1,600,000	Up to 250,000	8	4	1	15	15	16	50,000/burstable up to 110,000	3/burstable up to 6
ecs.r8y.4xlarge	16	128	10/25	3,000,000	400,000	25	8	1	30	30	16	80,000/burstable up to 110,000	5/burstable up to 6

ecs.r8y.8xlarge	32	256	16/none	5,000,000	750,000	32	8	1	30	30	16	125,000	8
ecs.r8y.16xlarge	64	512	32/none	10,000,000	1,500,000	32	8	1	30	30	32	250,000	16
ecs.r8y.32xlarge	128	1024	64/none	20,000,000	3,000,000	32	15	1	30	30	32	500,000	32

Note

To use the ecs.r8y.32xlarge instance type, submit a ticket.

g6r, general-purpose instance family

- **Introduction**: This instance family uses the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**: containers, microservices, scenarios where applications such as DevOps applications are developed and tested, websites, application servers, game servers, and CPU-based machine learning and inference.
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.8 GHz Ampere® Altra® processors to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
  - Provides high storage I/O performance based on large computing capacity.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.

g6r instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.g6r.large	2	8	1/10	900,000	Up to 250,000	2	3	6	1	12,500	1
ecs.g6r.xlarge	4	16	1.5/10	1,000,000	Up to 250,000	4	4	15	1	20,000	1.5
ecs.g6r.2xlarge	8	32	2.5/10	1,600,000	Up to 250,000	8	4	15	1	30,000	2
ecs.g6r.4xlarge	16	64	5/10	2,000,000	300,000	8	8	30	1	60,000	3
ecs.g6r.8xlarge	32	128	8/10	3,000,000	600,000	16	7	30	1	75,000	4
ecs.g6r.16xlarge	64	256	16/none	6,000,000	900,000	32	7	30	1	150,000	8

c6r, compute-optimized instance family

- **Introduction**: This instance family uses the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**:
  - Containers and microservices
  - Scenarios where applications such as DevOps applications are developed and tested
  - Websites and application servers
  - CPU-based machine learning and inference
  - High-performance nce and engineering applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.8 GHz Ampere® Altra® processors to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.

- Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
- Provides high storage I/O performance based on large computing capacity. For more information, see Storage I/O performance.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high packet forwarding rates.
  - Provides burstable network bandwidth for low-specification instances.
  - Provides high network performance based on large computing capacity.

The c6r instance family includes the following instance types: ecs.c6r.large, ecs.c6r.xlarge, ecs.c6r.2xlarge, ecs.c6r.4xlarge, ecs.c6r.8xlarge, and ecs.c6r.16xlarge. Click the following panel to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPU	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.c6r.large	2	4	1/10	900,000	Up to 250,000	2	3	6	1	12,500	1
ecs.c6r.xlarge	4	8	1.5/10	1,000,000	Up to 250,000	4	4	15	1	20,000	1.5
ecs.c6r.2xlarge	8	16	2.5/10	1,600,000	Up to 250,000	8	4	15	1	30,000	2
ecs.c6r.4xlarge	16	32	5/10	2,000,000	300,000	8	8	30	1	60,000	3
ecs.c6r.8xlarge	32	64	8/10	3,000,000	600,000	16	7	30	1	75,000	4
ecs.c6r.16xlarge	64	128	16/none	6,000,000	900,000	32	7	30	1	150,000	8

ECS Bare Metal Instance families

ebmgn8v, GPU-accelerated compute-optimized ECS Bare Metal Instance family

This instance family is available only in specific regions, including regions outside China. To use the instance family, contact Alibaba Cloud sales personnel.

- **Introduction**: This instance family is an 8th-generation GPU-accelerated compute-optimized ECS Bare Metal Instance family provided by Alibaba Cloud for AI model training and ultra-large models. Each instance of this instance family is equipped with eight GPUs.
- **Supported scenarios**:
  - Multi-GPU parallel inference computing for large language models (LLMs) that have more than 70 billion parameters
  - Traditional AI model training and autonomous driving training, for which each GPU delivers computing power of up to 39.5 TFLOPS in the single-precision floating-point format (FP32)
  - Small and medium-sized model training scenarios that leverage the NVLink connections among the eight GPUs
- **Benefits and positioning**:
  - **High-speed and large-capacity GPU memory**: Each GPU is equipped with 96 GB of HBM3E memory and delivers up to 4 TB/s of memory bandwidth, which greatly accelerates model training and inference.
  - **High bandwidth between GPUs**: Multiple GPUs are interconnected by using 900 GB/s NVLink connections. The efficiency of multi-GPU training and inference is much higher than that of previous generations of GPU-accelerated instances.
  - **Quantization of large models**: This instance family supports computing power in the 8-bit floating point format (FP8) and optimizes computing power for large-scale parameter training and inference. This significantly improves the computing speed of training and inference and reduces memory usage.
- **Compute**:
  - Uses the latest CIPU 1.0 processors.
    - Decouples computing capabilities from storage capabilities, allowing you to flexibly select storage resources based on your business requirements, and increases inter-instance bandwidth to 160 Gbit/s for faster data transmission and processing compared with 7th-generation instance families.
    - Uses the bare metal capabilities provided by CIPU processors to support peer-to-peer (P2P) communication between GPU-accelerated instances.
  - Uses the 4th-generation Intel Xeon Scalable processors that deliver an all-core turbo frequency of up to 3.1 GHz and provides 192 vCPUs.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, and elastic ephemeral disks (EEDs). For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high network performance with a packet forwarding rate of 30,000,000 pps.
  - Supports elastic RDMA interfaces (ERIs) to allow inter-instance RDMA-based communication in VPCs and provides up to 160 Gbit/s of bandwidth per instance, which is suitable for training tasks based on CV models and traditional models.

    > **Note**
    >
    > For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.

Instance types

Instance type	vCPUs	Memory (GiB)	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Private IPv4 addresses per ENI	IPv6 addresses per ENI	NIC queues (Primary ENI/Secondary ENI)	ENIs	Maximum attached data disks	Maximum disk bandwidth (Gbit/s)
ecs.ebmgn8v.48xlarge	192	1,024	96GB*8	160 (80 × 2)	30,000,000	30	30	64	32	31	6

> **Note**
>
> The boot mode of the images that are used by instances of this instance family must be UEFI. If you want to use custom images on the instances, make sure that the images support the UEFI boot mode and the boot mode of the images is set to UEFI. For information about how to set the boot mode of a custom image, see Set the boot mode of custom images to the UEFI mode by calling API operations.

ebmgn8is, GPU-accelerated compute-optimized ECS Bare Metal Instance family

This instance family is available only in specific regions, including regions outside China. To use the instance family, contact Alibaba Cloud sales personnel.

- **Introduction**: This instance family is an 8th-generation GPU-accelerated compute-optimized ECS Bare Metal instance family provided by Alibaba Cloud in response to the recent developments in the AI generation field. Each instance of this instance family is equipped with eight GPUs.
- **Supported scenarios:**
  - Production and rendering of special effects for animation, film, and television based on workstation-level graphics processing capabilities in scenarios in which Alibaba Cloud Marketplace GRID images are used, the GRID driver is installed, and OpenGL and Direct3D graphics capabilities are enabled
  - Scenarios in which the management services provided by Container Service for Kubernetes (ACK) for containerized applications are used to support AI-generated graphic content and LLM inference tasks with up to 130 billion parameters
  - Other general-purpose AI recognition, image recognition, and speech recognition scenarios
- **Benefits and positioning:**
  - **Graphic processing**: This instance family uses high-frequency 5th-generation Intel Xeon Scalable processors to deliver sufficient CPU computing power in 3D modeling scenarios and achieve smooth graphics rendering and design.
  - **Inference tasks**: This instance family uses innovative GPUs, each with 48 GB of memory, which accelerate inference tasks and support the FP8 floating-point format. You can use this instance family together with ACK to support the inference of various AI-generated content (AIGC) models and accommodate inference tasks for LLMs that have less than 70 billion parameters.
  - **Training tasks**: This instance family provides cost-effective computing capabilities and delivers the FP32 computing performance double that of the 7th-generation inference instances. Instances of this instance family are suitable for training FP32-based CV models and other small and medium-sized models.

- **Uses the latest CIPU 1.0 processors that provide the following benefits:**
  - Decouples computing capabilities from storage capabilities, allowing you to flexibly select storage resources based on your business requirements, and increases inter-instance bandwidth to 160 Gbit/s for faster data transmission and processing compared with previous-generation instance families.
  - Uses the bare metal capabilities provided by CIPU processors to support Peripheral Component Interconnect Express (PCIe) P2P communication between GPU-accelerated instances.
- **Compute:**
  - Uses innovative GPUs that have the following features:
    - Support for acceleration features such as vGPU, RTX technology, and TensorRT inference engine
    - Support for PCIe Switch interconnect, which achieves a 36% increase in NVIDIA Collective Communications Library (NCCL) performance compared with the CPU direct connection scheme and helps improve inference performance by up to 9% when you run LLM inference tasks on multiple GPUs in parallel
    - Support for eight GPUs per instance with 48 GB of memory per GPU to support LLM inference tasks with 70 billion or more parameters on a single instance
  - Uses 3.4 GHz Intel® Xeon® Scalable (SPR) processors that deliver an all-core turbo frequency of up to 3.9 GHz.
- **Storage:**
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, and EEDs. For information about disks, see Overview of Block Storage.
- **Network:**
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high network performance with a packet forwarding rate of 30,000,000 pps.
  - Supports ERIs to allow inter-instance RDMA-based communication in VPCs and provides up to 160 Gbit/s of bandwidth per instance, which is suitable for training tasks based on CV models and traditional models.

    > **Note**
    >
    > For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.

Instance types

Instance type	vCPUs	Memory (GiB)	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Private IPv4 addresses per ENI	IPv6 addresses per ENI	NIC queues (Primary ENI/Secondary ENI)	ENIs	Maximum attached data disks	Maximum disk bandwidth (Gbit/s)
ecs.ebmgn8is.32xlarge	128	1,024	48 GB × 8	160 (80 × 2)	30,000,000	30	30	64/16	32	31	6

> **Note**
>
> The boot mode of the images that are used by instances of this instance family must be UEFI. If you want to use custom images on the instances, make sure that the images support the UEFI boot mode and the boot mode of the images is set to UEFI. For information about how to set the boot mode of a custom image, see Set the boot mode of custom images to the UEFI mode by calling API operations.

ebmgn7e, GPU-accelerated compute-optimized ECS Bare Metal Instance family

- **Introduction**: This instance family uses the SHENLONG architecture to provide flexible and powerful software-defined compute.
- **Supported scenarios:**
  - Deep learning training and development
  - High-performance computing (HPC) and simulations

> **Important**
>
> When you use AI training services that feature a high communication load, such as transformer models, you must enable NVLink for GPU-to-GPU communication. Otherwise, data may be damaged due to unpredictable failures that are caused by large-scale data transmission over Peripheral Component Interconnect Express (PCIe) links. If you do not understand the topology of the communication links that are used for AI training services, submit a ticket to obtain technical support.

- **Compute:**
  - Uses 2.9 GHz Intel® Xeon® Scalable processors that deliver an all-core turbo frequency of 3.5 GHz and supports PCIe 4.0 interfaces.
- **Storage:**
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- **Network:**
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues (Primary NIC/Secondary NIC)	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.ebmgn7e.32xlarge	128	1,024	80 GB × 8	64	24,000,000	32/12	32	10	1

You must check the status of the multi-instance GPU (MIG) feature and enable or disable the MIG feature after you start an ebmgn7e instance. For information about the MIG feature, see NVIDIA Multi-Instance GPU User Guide.

The following table describes whether the MIG feature is supported by the instance types in the ebmgn7e instance family.

Instance type	Support for MIG	Description
ecs.ebmgn7e.32xlarge	Yes	The MIG feature is supported by ebmgn7e instances.

ebmgn7i, GPU-accelerated compute-optimized ECS Bare Metal Instance family

- **Introduction**: This instance family uses the SHENLONG architecture to provide flexible and powerful software-defined compute.
- **Supported scenarios:**
  - Concurrent AI inference tasks that require high-performance CPUs, memory, and GPUs, such as image recognition, speech recognition, and behavior identification
  - Compute-intensive graphics processing tasks that require high-performance 3D graphics virtualization capabilities, such as remote graphic design and cloud gaming
  - Scenarios that require high network bandwidth and disk bandwidth, such as the creation of high-performance render farms
  - Small-scale deep learning and training applications that require high network bandwidth
- **Compute:**
  - Uses NVIDIA A10 GPUs that have the following features:
    - Innovative NVIDIA Ampere architecture
    - Support for acceleration features such as vGPU, RTX technology, and TensorRT inference engine
  - Uses 2.9 GHz Intel® Xeon® Scalable (Ice Lake) processors that deliver an all-core turbo frequency of 3.5 GHz.
- **Storage:**
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- **Network:**
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	GPU	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.ebmgn7i.32xlarge	128	768	NVIDIA A10 × 4	24 GB × 4	64	24,000,000	32	32	10	1

ebmgn7, GPU-accelerated compute-optimized ECS Bare Metal Instance family

- **Introduction**: This instance family uses the SHENLONG architecture to provide flexible and powerful software-defined compute.
- **Supported scenarios:**
  - Deep learning applications, such as training applications of AI algorithms used in image classification, autonomous vehicles, and speech recognition
  - Scientific computing applications that require robust GPU computing capabilities, such as computational fluid dynamics, computational finance, molecular dynamics, and environmental analytics
- **Compute:**
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors.
- **Storage:**
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- **Network:**
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

Instance types

Instance type	vCPUs	Memory (GiB)	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.ebmgn7.26xlarge	104	768	40 GB × 8	30	18,000,000	16	15	10	1

ebmgn6ia, GPU-accelerated compute-optimized ECS Bare Metal Instance family

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.
  - This instance family uses NVIDIA T4 GPUs to offer GPU acceleration capabilities for graphics and AI applications and adopts container technology to start at least 60 virtual Android devices and provide hardware-accelerated video transcoding.
- **Supported scenarios:**
  - Remote application services based on Android, such as always-on cloud-based services, cloud-based mobile games, cloud-based mobile phones, and Android service crawlers.
- **Compute:**
  - Offers a CPU-to-memory ratio of 1:3.
  - Uses 2.8 GHz Ampere® Altra® Arm-based processors that deliver a turbo frequency of 3.0 GHz and provides high performance and high compatibility with applications for Android servers.

- **Storage:**
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- **Network:**
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.

Instance types

Instance type	vCPUs	Memory (GiB)	GPU	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.ebmgn6ia.20xlarge	80	256	NVIDIA T4 × 2	16 GB × 2	32	24,000,000	32	15	10	1

> **Note**
>
> Ampere® Altra® processors have specific requirements for operating system kernels. Instances of the preceding instance type can use Alibaba Cloud Linux 3 images and CentOS 8.4 or later images. We recommend that you use Alibaba Cloud Linux 3 images on the instances. If you want to use another operating system distribution, patch the kernel of an instance that runs an operating system of that distribution, create a custom image from the instance, and then use the custom image to create instances of the instance type. For information about kernel patches, visit Ampere Altra (TM) Linux Kernel Porting Guide.

ebmgn6e, GPU-accelerated compute-optimized ECS Bare Metal Instance family

- **Introduction:**
  - This instance family uses the SHENLONG architecture to provide flexible and powerful software-defined compute.
  - This instance family uses NVIDIA V100 GPUs that each have 32 GB of GPU memory and support NVLink.
  - This instance family uses NVIDIA V100 GPUs (SXM2-based) that have the following features:
    - Innovative NVIDIA Volta architecture
    - 32 GB of HBM2 memory (900 GB/s bandwidth) per GPU
    - 5,120 CUDA cores per GPU
    - 640 Tensor cores per GPU
    - Up to six NVLink connections per GPU, each of which provides a bandwidth of 25 GB/s in each direction for a total bandwidth of 300 GB/s (6 × 25 × 2 = 300)
- **Supported scenarios:**
  - Deep learning applications, such as training and inference applications of AI algorithms used in image classification, autonomous vehicles, and speech recognition
  - Scientific computing applications, such as computational fluid dynamics, computational finance, molecular dynamics, and environmental analytics
- **Compute:**

- Offers a CPU-to-memory ratio of 1:8.
- Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors.

- **Storage:**
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
- **Network:**
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

Instance types

Instance type	vCPUs	Memory (GiB)	GPU	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.ebmgn6e.24xlarge	96	768	NVIDIA V100 × 8	32 GB × 8	32	4,800,000	16	15	10	1

ebmgn6v, GPU-accelerated compute-optimized ECS Bare Metal Instance family

- **Introduction:**
  - This instance family uses the SHENLONG architecture to provide flexible and powerful software-defined compute.
  - This instance family uses NVIDIA V100 GPUs.
  - This instance family uses NVIDIA V100 GPUs (SXM2-based) that have the following features:
    - Innovative NVIDIA Volta architecture
    - 16 GB of HBM2 memory (900 GB/s bandwidth) per GPU
    - 5,120 CUDA cores per GPU
    - 640 Tensor cores per GPU
    - Up to six NVLink connections per GPU, each of which provides a bandwidth of 25 GB/s in each direction for a total bandwidth of 300 GB/s (6 × 25 × 2 = 300)
- **Supported scenarios:**
  - Deep learning applications, such as training and inference applications of AI algorithms used in image classification, autonomous vehicles, and speech recognition
  - Scientific computing applications, such as computational fluid dynamics, computational finance, molecular dynamics, and environmental analytics
- **Compute:**
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors.
- **Storage:**
  - Is an instance family in which all instances are I/O optimized.

- Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.

- **Network:**
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

Instance types

Instance type	vCPUs	Memory (GiB)	GPU	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.ebmgn6v.24xlarge	96	384	NVIDIA V100 × 8	16 GB × 8	30	4,500,000	8	32	10	1

ebmgn6i, GPU-accelerated compute-optimized ECS Bare Metal Instance family

- **Introduction:**
  - This instance family uses the SHENLONG architecture to provide flexible and powerful software-defined compute.
  - This instance family uses NVIDIA T4 GPUs that have the following features:
    - Innovative NVIDIA Turing architecture
    - 16 GB of memory (320 GB/s bandwidth) per GPU
    - 2,560 CUDA cores per GPU
    - Up to 320 Turing Tensor cores per GPU
    - Mixed-precision Tensor cores that support 65 FP16 TFLOPS, 130 INT8 TOPS, and 260 INT4 TOPS

- **Supported scenarios:**
  - AI (deep learning and machine learning) inference for computer vision, voice recognition, speech synthesis, natural language processing (NLP), machine translation, and reference systems
  - Real-time rendering for cloud gaming
  - Real-time rendering for Augmented Reality (AR) and Virtual Reality (VR) applications
  - Graphics workstations or graphics-heavy computing
  - GPU-accelerated databases
  - High-performance computing

- **Compute:**
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors.

- **Storage:**
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.

- **Network:**

- Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
- Provides high network performance based on large computing capacity.

Instance types

Instance type	vCPUs	Memory (GiB)	GPU	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.ebmgn6i.24xlarge	96	384	NVIDIA T4 × 4	16 GB × 4	30	4,500,000	8	32	10	1

ebmc8i, compute-optimized ECS Bare Metal Instance family

- **Introduction**: This instance family uses the innovative CIPU architecture developed by Alibaba Cloud to provide stable computing power, a more robust I/O engine, and dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Web frontend servers
  - Frontend servers of massively multiplayer online (MMO) games
  - Data analytics, batch processing, and video encoding
  - High-performance scientific and engineering applications
- **Compute**:
  - Uses Intel® Xeon® Emerald Rapids or Intel® Xeon® Sapphire Rapids processors that deliver a clock speed of at least 2.7 GHz and an all-core turbo frequency of 3.2 GHz to provide consistent computing performance.

    > **Note**
    >
    > When you purchase an instance of this instance family, the system randomly allocates one type of the preceding processors to the instance. You cannot select a processor type for the instance.

  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see Operating system versions that support AMD Genoa processors used by eighth-generation AMD instance types.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports the Non-Volatile Memory Express (NVMe) protocol. For more information, see NVMe protocol.

- Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports elastic RDMA interfaces (ERIs). For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.ebmc8i.48xlarge	192	512	100/none	30,000,000	4,000,000	64 (primary ENI)/16 (secondary ENI)	72	30	30	1,000,000/none	48/none

ebmc7, compute-optimized ECS Bare Metal Instance family

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.
  - This instance family provides dedicated hardware resources and physical isolation.

- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Web frontend servers
  - Frontend servers of MMO games
  - Data analytics, batch processing, and video encoding
  - High-performance scientific and engineering applications

- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.9 GHz Intel ® Xeon ® Platinum 8369B (Ice Lake) processors that deliver an all-core turbo frequency of 3.5 GHz.

- **Storage**:

- Is an instance family in which all instances are I/O optimized.
- Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmc7.32xlarge	128	256	64	24,000,000	2,400,000	32	20	20	600,000	32

ebmc7a, compute-optimized ECS Bare Metal Instance family

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.
  - This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Video encoding, decoding, and rendering
  - Data analytics and computing
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.55 GHz AMD EPYC™ MILAN processors that deliver a single-core turbo frequency of up to 3.5 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmc7a.64xlarge	256	512	64	24,000,000	4,000,000	32	31	15	1	600,000	32

Note

- The boot mode of the images that are used by instances of this instance family must be UEFI. If you want to use custom images on the instances, make sure that the boot mode of the images is set to UEFI. For information about how to set the boot mode of a custom image, see Set the boot mode of custom images to the UEFI mode by calling API operations.
- Ubuntu 18 and Debian 9 operating system kernels do not support AMD EPYCTM MILAN processors. Do not use Ubuntu 18 or Debian 9 images to create instances of this instance family. Instances of this instance family that are created from Ubuntu 18 or Debian 9 images cannot start.

ebmc6me, compute-optimized ECS Bare Metal Instance family

- **Introduction**: This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Video encoding, decoding, and rendering
  - Frontend servers of MMO games
  - High-performance scientific and engineering applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:3.
  - Uses 2.3 GHz Intel® Xeon® Gold 5218 (Cascade Lake) processors that deliver a turbo frequency of 3.9 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance with a packet forwarding rate of 6,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmc6me.16xlarge	64	192	32	6,000,000	1,800,000	32	10	1	200,000	16

ebmc6a, compute-optimized ECS Bare Metal Instance family

This instance family is in invitational preview. To use the instance family, submit a ticket.

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.
  - This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Video encoding, decoding, and rendering
  - Data analytics and computing
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.6 GHz AMD EPYC™ ROME processors that deliver a turbo frequency of 3.3 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)

ecs.ebmc6a.64xlarge	256	512	64	24,000,000	32	31	10	1	600,000	32

Note

The boot mode of the images that are used by instances of this instance family must be UEFI. If you want to use custom images on the instances, make sure that the boot mode of the images is set to UEFI. For information about how to set the boot mode of a custom image, see Set the boot mode of custom images to the UEFI mode by calling API operations.

ebmc6e, performance-enhanced compute-optimized ECS Bare Metal Instance family

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.
  - This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Web frontend servers
  - Frontend servers of MMO games
  - Data analytics, batch processing, and video encoding
  - High-performance scientific and engineering applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver an all-core turbo frequency of 3.2 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmc6e.26xlarge	104	192	32	24,000,000	1,800,000	32	10	1	480,000	16

ebmc6, compute-optimized ECS Bare Metal Instance family

- **Introduction**: This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Video encoding, decoding, and rendering
  - Frontend servers of MMO games
  - High-performance scientific and engineering applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver an all-core turbo frequency of 3.2 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance with a packet forwarding rate of 6,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmc6.26xlarge	104	192	32	6,000,000	1,800,000	32	20	1	200,000	16

ebmg8i, general-purpose ECS Bare Metal Instance family

- **Introduction**: This instance family uses the innovative Cloud Infrastructure Processing Unit (CIPU) architecture developed by Alibaba Cloud to provide stable computing power, a more robust I/O engine, and dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Enterprise-level applications of various types and sizes
  - Websites and application servers
  - Game servers
  - Small and medium-sized database systems, caches, and search clusters
  - Data analytics and computing
  - High-performance scientific and engineering applications
- **Compute**:
  - Uses Intel® Xeon® Emerald Rapids or Intel® Xeon® Sapphire Rapids processors that deliver a clock speed of at least 2.7 GHz and an all-core turbo frequency of 3.2 GHz to provide consistent computing performance.

    > **Note**
    >
    > When you purchase an instance of this instance family, the system randomly allocates one type of the preceding processors to the instance. You cannot select a processor type for the instance.

  - Supports Hyper-Threading. By default, Hyper-Threading is enabled. For more information, see Specify and view CPU options.
  - Is compatible with specific operating systems. For more information, see Compatibility between Intel instance types and operating systems.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports the Non-Volatile Memory Express (NVMe) protocol. For more information, see NVMe protocol.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports elastic RDMA interfaces (ERIs). For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline/burst IOPS	Disk baseline/burst bandwidth (Gbit/s)
ecs.ebmg8i.48xlarge	192	1,024	100/none	30,000,000	4,000,000	64 (primary ENI)/16 (secondary ENI)	72	30	30	1,000,000/none	48/none

ebmg7, general-purpose ECS Bare Metal Instance family

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.
  - This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Enterprise-level applications of various types and sizes
  - Websites and application servers
  - Game servers
  - Small and medium-sized database systems, caches, and search clusters
  - Data analytics and computing
  - High-performance scientific and engineering applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.9 GHz Intel® Xeon® Platinum 8369B (Ice Lake) processors that deliver an all-core turbo frequency of 3.5 GHz.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmg7.32xlarge	128	512	64	24,000,000	2,400,000	32	20	20	600,000	32

ebmg7a, general-purpose ECS Bare Metal Instance family

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.
  - This instance family provides dedicated hardware resources and physical isolation.
- Supported scenarios:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Computing clusters and memory-intensive data processing
  - Video encoding, decoding, and rendering
  - Data analytics and computing
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.55 GHz AMD EPYC™ MILAN processors that deliver a single-core turbo frequency of up to 3.5 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.

- Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmg7a.64xlarge	256	1,024	64	24,000,000	4,000,000	32	31	15	1	600,000	32

Note

- The boot mode of the images that are used by instances of this instance family must be Unified Extensible Firmware Interface (UEFI). If you want to use custom images on the instances, make sure that the boot mode of the images is set to UEFI. For information about how to set the boot mode of a custom image, see Set the boot mode of custom images to the UEFI mode by calling API operations.
- Ubuntu 18 and Debian 9 operating system kernels do not support AMD EPYC™ MILAN processors. Do not use Ubuntu 18 or Debian 9 images to create instances of this instance family. Instances of this instance family that are created from Ubuntu 18 or Debian 9 images cannot start.

ebmg6a, general-purpose ECS Bare Metal Instance family

This instance family is in invitational preview. To use this instance family, submit a ticket.

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.
  - This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Video encoding, decoding, and rendering
  - Computing clusters and memory-intensive data processing
  - Data analytics and computing
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.6 GHz AMD EPYC™ ROME processors that deliver a turbo frequency of 3.3 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.

- Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmg6a.64xlarge	256	1,024	64	24,000,000	32	31	10	1	600,000	32

> **Note**
>
> The boot mode of the images that are used by instances of this instance family must be UEFI. If you want to use custom images on the instances, make sure that the boot mode of the images is set to UEFI. For information about how to set the boot mode of a custom image, see Set the boot mode of custom images to the UEFI mode by calling API operations.

ebmg6e, performance-enhanced general-purpose ECS Bare Metal Instance family

Features:

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.
  - This instance family provides dedicated hardware resources and physical isolation.

- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Enterprise-level applications of various types and sizes
  - Websites and application servers
  - Game servers
  - Small and medium-sized database systems, caches, and search clusters
  - Data analytics and computing
  - Computing clusters and memory-intensive data processing
  - High-performance scientific and engineering applications

- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver an all-core turbo frequency of 3.2 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmg6e.26xlarge	104	384	32	24,000,000	1,800,000	32	10	1	480,000	16

ebmg6, general-purpose ECS Bare Metal Instance family

Features:

- **Introduction**: This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Video encoding, decoding, and rendering
  - Enterprise-level applications such as large and medium-sized databases
  - Computing clusters and memory-intensive data processing
  - Data analytics and computing
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver an all-core turbo frequency of 3.2 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance with a packet forwarding rate of 6,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmg6.26xlarge	104	384	32	6,000,000	1,800,000	32	20	1	200,000	16

> **Note**
>
> The CPU monitoring information about ECS bare metal instances cannot be obtained. To obtain the CPU monitoring information about an ECS bare metal instance, install the CloudMonitor agent on the instance. For more information, see Install and uninstall the CloudMonitor agent.

ebmr7, memory-optimized ECS Bare Metal Instance family

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.
  - This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - High-performance databases and in-memory databases
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters
  - High-performance scientific and engineering applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses 2.9 GHz Intel ® Xeon ® Platinum 8369B (Ice Lake) processors that deliver an all-core turbo frequency of 3.5 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- **Network**:

- Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
- Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmr7.32xlarge	128	1,024	64	24,000,000	2,400,000	32	20	20	600,000	32

ebmr7a, memory-optimized ECS Bare Metal Instance family

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.
  - This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - In-memory databases
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses 2.55 GHz AMD EPYC™ MILAN processors that deliver a maximum single-core turbo frequency of 3.5 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmr7a.64xlarge	256	2,048	64	24,000,000	4,000,000	32	31	15	1	600,000	32

Note

- The boot mode of the images that are used by instances of this instance family must be UEFI. If you want to use custom images on the instances, make sure that the boot mode of the images is set to UEFI. For information about how to set the boot mode of a custom image, see Set the boot mode of custom images to the UEFI mode by calling API operations.
- Ubuntu 18 and Debian 9 operating system kernels do not support AMD EPYC™ MILAN processors. Do not use Ubuntu 18 or Debian 9 images to create instances of this instance family. Instances of this instance family that are created from Ubuntu 18 or Debian 9 images cannot start.

ebmr6a, memory-optimized ECS Bare Metal Instance family

This instance family is in invitational preview. To use the instance family, submit a ticket.

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.
  - This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - In-memory databases
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses 2.6 GHz AMD EPYC™ ROME processors that deliver a turbo frequency of 3.3 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.

- Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmr6a.64xlarge	256	2,048	64	24,000,000	32	31	10	1	600,000	32

> **Note**
>
> The boot mode of the images that are used by instances of this instance family must be UEFI. If you want to use custom images on the instances, make sure that the boot mode of the images is set to UEFI. For information about how to set the boot mode of a custom image, see Set the boot mode of custom images to the UEFI mode by calling API operations.

ebmr6e, performance-enhanced memory-optimized ECS Bare Metal Instance family

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.
  - This instance family provides dedicated hardware resources and physical isolation.

- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - High-performance databases and in-memory databases
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters
  - High-performance scientific and engineering applications

- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver an all-core turbo frequency of 3.2 GHz.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmr6e.26xlarge	104	768	32	24,000,000	1,800,000	32	10	1	480,000	16

ebmr6, memory-optimized ECS Bare Metal Instance family

- **Introduction**: This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - High-performance databases and in-memory databases
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver an all-core turbo frequency of 3.2 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance with a packet forwarding rate of 6,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmr6.26xlarge	104	768	32	6,000,000	1,800,000	32	20	1	200,000	16

ebmre6p, persistent memory-optimized ECS Bare Metal Instance family

To use the ebmre6p instance family, submit a ticket.

- **Introduction**: This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - In-memory databases such as Redis
  - High-performance databases such as SAP HANA
  - Other memory-intensive applications such as AI applications and smart search applications
- **Compute**:
  - Uses the Intel® Optane™ persistent memory and is tuned for Redis applications in an end-to-end manner to provide cost-effectiveness.
  - Supports a total memory capacity of up to 1,920 GiB (384 GiB of DRAM + 1,536 GiB of Intel® Optane™ persistent memory), offers a CPU-to-memory ratio of 1:20, and can meet the needs of memory-intensive applications.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver an all-core turbo frequency of 3.2 GHz to provide consistent computing performance.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance with a packet forwarding rate of 6,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Persistent memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)

ecs.ebmre6p.26xlarge	104	384	1,536	32	6,000,000	32	10	1	200,000	16

ebmre6-6t, performance-enhanced memory-optimized ECS Bare Metal Instance family

To use the ebmre6-6t instance family, submit a ticket.

- **Introduction**: This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - In-memory databases and high-performance databases such as SAP HANA
  - Memory-intensive applications
  - Big data processing engines such as Apache Spark and Presto
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:30.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8269 (Cascade Lake) processors that deliver an all-core turbo frequency of 3.2 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance with a packet forwarding rate of 6,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmre6-6t.52xlarge	208	6,144	32	6,000,000	1,800,000	32	10	1	200,000	16

ebmhfg7, general-purpose ECS Bare Metal Instance family with high clock speeds

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.

- This instance family provides dedicated hardware resources and physical isolation.

- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Enterprise-level applications of various types and sizes
  - Game servers
  - Small and medium-sized database systems, caches, and search clusters
  - High-performance scientific computing
  - Video encoding applications
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses third-generation Intel® Xeon® Scalable (Cooper Lake) processors that deliver a base frequency of at least 3.3 GHz and an all-core turbo frequency of 3.8 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmhfg7.48xlarge	192	768	64	24,000,000	32	31	10	1	600,000	32

ebmhfc7, compute-optimized ECS Bare Metal Instance family with high clock speeds

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.
  - This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware

- Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
- High-performance frontend server clusters
- Frontend servers of MMO games
- Data analytics, batch processing, and video encoding
- High-performance scientific and engineering applications

- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses third-generation Intel® Xeon® Scalable (Cooper Lake) processors that deliver a base frequency of at least 3.3 GHz and an all-core turbo frequency of 3.8 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports only ESSDs and ESSD AutoPL disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmhfc7.48xlarge	192	384	64	24,000,000	32	31	10	1	600,000	32

ebmhfr7, memory-optimized ECS Bare Metal Instance family with high clock speeds

- **Introduction:**
  - This instance family uses the third-generation SHENLONG architecture and fast path acceleration on chips to provide predictable and consistent ultra-high computing, storage, and network performance.
  - This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - High-performance databases and in-memory databases
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.

- Uses third-generation Intel® Xeon® Scalable (Cooper Lake) processors that deliver a base frequency of at least 3.3 GHz and an all-core turbo frequency of 3.8 GHz.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks. For information about disks, see Overview of Block Storage.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides ultra-high network performance with a packet forwarding rate of 24,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmhfr7.48xlarge	192	1,536	64	24,000,000	32	31	10	1	600,000	32

ebmhfg6, general-purpose ECS Bare Metal Instance family with high clock speeds

- **Introduction**: This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Enterprise-level applications such as large and medium-sized databases
  - Video encoding, decoding, and rendering
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.8.
  - Uses 3.1 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver an all-core turbo frequency of 3.5 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance with a packet forwarding rate of 6,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmhfg6.20xlarge	80	384	32	6,000,000	1,800,000	32	20	1	200,000	16

ebmhfc6, compute-optimized ECS Bare Metal Instance family with high clock speeds

- **Introduction**: This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Video encoding, decoding, and rendering
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.4.
  - Uses 3.1 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver an all-core turbo frequency of 3.5 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance with a packet forwarding rate of 6,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmhfc6.20xlarge	80	192	32	6,000,000	1,800,000	32	20	1	200,000	16

ebmhfr6, memory-optimized ECS Bare Metal Instance family with high clock speeds

- **Introduction**: This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - High-performance databases and in-memory databases
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:9.6.
  - Uses 3.1 GHz Intel® Xeon® Platinum 8269CY (Cascade Lake) processors that deliver an all-core turbo frequency of 3.5 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance with a packet forwarding rate of 6,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmhfr6.20xlarge	80	768	32	6,000,000	1,800,000	32	20	1	200,000	16

High-performance computing and SCC instance families

scchfc6, compute-optimized SCC instance family with high clock speeds

To use it, submit a ticket.

- **Introduction**: This instance family provides all features of ECS Bare Metal Instance. For more information, see Overview of ECS Bare Metal Instance families.
- **Supported scenarios**:
  - Large-scale machine learning training
  - Large-scale high-performance scientific computing and simulations
  - Large-scale data analytics, batch processing, and video encoding
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.4.
  - Uses 3.1 GHz Intel® Xeon® Platinum 8269 (Cascade Lake) processors that deliver an all-core turbo frequency of 3.5 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports enhanced SSDs (ESSDs), ESSD AutoPL disks, standard SSDs, and ultra disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports both RoCE networks and VPCs. RoCE networks are dedicated to RDMA communication.

Instance types

Instance type	vCPU	Physical cores	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	RoCE network bandwidth (Gbit/s)	ENIs
ecs.scchfc6.20xlarge	80	40	192.0	30	6,000,000	50	32

> **Note**
>
> ecs.scchfc6.20xlarge provides 80 logical processors on 40 physical cores.

scchfg6, general-purpose SCC instance family with high clock speeds

To use it, submit a ticket.

- **Introduction**: This instance family provides all features of ECS Bare Metal Instance. For more information, see Overview of ECS Bare Metal Instance families.
- **Supported scenarios**:
  - Large-scale machine learning training
  - Large-scale high-performance scientific computing and simulations
  - Large-scale data analytics, batch processing, and video encoding
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.8.
  - Uses 3.1 GHz Intel® Xeon® Platinum 8269 (Cascade Lake) processors that deliver an all-core turbo frequency of 3.5 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.

- Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports both RoCE networks and VPCs. RoCE networks are dedicated to RDMA communication.

Instance types

Instance type	vCPU	Physical cores	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	RoCE network bandwidth (Gbit/s)	ENIs
ecs.scchfg6.20xlarge	80	40	384.0	30	6,000,000	50	32

> **Note**
>
> ecs.scchfg6.20xlarge provides 80 logical processors on 40 physical cores.

scchfr6, memory-optimized SCC instance family with high clock speeds

To use it, submit a ticket.

- **Introduction**: This instance family provides all features of ECS Bare Metal Instance. For more information, see Overview of ECS Bare Metal Instance families.
- **Supported scenarios**:
  - Large-scale machine learning training
  - Large-scale high-performance scientific computing and simulations
  - Large-scale data analytics, batch processing, and video encoding
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:9.6.
  - Uses 3.1 GHz Intel® Xeon® Platinum 8269 (Cascade Lake) processors that deliver an all-core turbo frequency of 3.5 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports both RoCE networks and VPCs. RoCE networks are dedicated to RDMA communication.

Instance types

Instance type	vCPU	Physical cores	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	RoCE network bandwidth (Gbit/s)	ENIs
ecs.scchfr6.20xlarge	80	40	768.0	30	6,000,000	50	32

> **Note**
>
> ecs.scchfr6.20xlarge provides 80 logical processors on 40 physical cores.

scch5, SCC instance family with high clock speeds

- **Introduction**: This instance family provides all features of ECS Bare Metal Instance. For more information, see Overview of ECS Bare Metal Instance families.
- **Supported scenarios**:
  - Large-scale machine learning training
  - Large-scale high-performance scientific computing and simulations
  - Large-scale data analytics, batch processing, and video encoding
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:3.
  - Uses 3.1 GHz Intel® Xeon® Gold 6149 (Skylake) processors.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports only standard SSDs and ultra disks.
- **Network**:
  - Supports only IPv4.
  - Supports both RoCE networks and VPCs. RoCE networks are dedicated to RDMA communication.

Instance types

Instance type	vCPU	Physical cores	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	RoCE network bandwidth (Gbit/s)	ENIs
ecs.scch5.16xlarge	64	32	192.0	10	4,500,000	50	32

> **Note**
>
> ecs.scch5.16xlarge provides 64 logical processors on 32 physical cores.

ebmc5s, network-enhanced compute-optimized ECS Bare Metal Instance family

- **Introduction**: This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Scenarios where large volumes of packets are received and transmitted, such as live commenting on videos and telecom data forwarding
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch

- Video encoding, decoding, and rendering

- **Compute**:
  - Offers a CPU-to-memory ratio of 1:2.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors that deliver an all-core turbo frequency of 2.7 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance with a packet forwarding rate of 4,500,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmc5s.24xlarge	96	192	32	4,500,000	1,800,000	32	10	1	200,000	16

ebmg5s, network-enhanced general-purpose ECS Bare Metal Instance family

- **Introduction**: This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Enterprise-level applications such as large and medium-sized databases
  - Video encoding
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors that deliver an all-core turbo frequency of 2.7 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.

- Provides high network performance with a packet forwarding rate of 4,500,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.ebmg5s.24xlarge	96	384	32	4,500,000	1,800,000	32	10	1	200,000	16

ebmr5s, network-enhanced memory-optimized ECS Bare Metal Instance family

- **Introduction**: This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - High-performance databases and in-memory databases
  - Data analytics, data mining, and distributed memory caching
  - Enterprise-level memory-intensive applications such as Hadoop clusters and Spark clusters
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:8.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors that deliver an all-core turbo frequency of 2.7 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports only IPv4.
  - Provides high network performance with a packet forwarding rate of 4,500,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Connections	ENIs	Private IPv4 addresses per ENI	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)

ecs.ebmr5s.24xlarge	96	768	32	4,500,000	1,800,000	32	10	200,000	16

ebmg5, general-purpose ECS Bare Metal Instance family

- **Introduction**: This instance family provides dedicated hardware resources and physical isolation.
- **Supported scenarios**:
  - Workloads that require direct access to physical resources or that require a license to be bound to the hardware
  - Scenarios that require compatibility with third-party hypervisors to implement hybrid-cloud and multi-cloud deployments
  - Containers such as Docker, Clear Containers, and Pouch
  - Enterprise-level applications such as large and medium-sized databases
  - Video encoding
- **Compute**:
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors that deliver an all-core turbo frequency of 2.7 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks. For information about disks, see Overview of Block Storage.
- **Network**:
  - Supports only IPv4.
  - Provides high network performance with a packet forwarding rate of 4,000,000 pps.

Instance types

Instance type	vCPUs	Memory (GiB)	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	ENIs	Private IPv4 addresses per ENI
ecs.ebmg5.24xlarge	96	384	10	4,000,000	32	10

Enterprise-level heterogeneous computing instance families

sgn7i-vws, vGPU-accelerated instance family with shared CPUs

- **Introduction**:

- This instance family uses the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude. This way, data storage and model loading can be performed more quickly.
- Instances of this instance family share CPU and network resources to maximize the utilization of underlying resources. Each instance has exclusive access to its memory and GPU memory to provide data isolation and performance assurance.

  > **Note**
  >
  > If you want to use exclusive CPU resources, select the vgn7i-vws instance family.

- This instance family comes with an NVIDIA GRID vWS license and provides certified graphics acceleration capabilities for Computer Aided Design (CAD) software to meet the requirements of professional graphic design. Instances of this instance family can serve as lightweight GPU-accelerated compute-optimized instances to reduce the costs of small-scale AI inference tasks.

- **Supported scenarios**:
  - Concurrent AI inference tasks that require high-performance CPUs, memory, and GPUs, such as image recognition, speech recognition, and behavior identification
  - Compute-intensive graphics processing tasks that require high-performance 3D graphics virtualization capabilities, such as remote graphic design and cloud gaming
  - 3D modeling in fields that require the use of Ice Lake processors, such as animation and film production, cloud gaming, and mechanical design
- **Compute**:
  - Uses NVIDIA A10 GPUs that have the following features:
    - Innovative NVIDIA Ampere architecture
    - Support for acceleration features, such as vGPU, RTX, and TensorRT, to provide diversified business support
  - Uses 2.9 GHz Intel® Xeon® Scalable (Ice Lake) processors that deliver an all-core turbo frequency of 3.5 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports Enterprise SSDs (ESSDs) and ESSD AutoPL disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

The sgn7i-vws instance family includes the following instance types: ecs.sgn7i-vws-m2.xlarge, ecs.sgn7i-vws-m4.2xlarge, ecs.sgn7i-vws-m8.4xlarge, ecs.sgn7i-vws-m2s.xlarge, ecs.sgn7i-vws-m4s.2xlarge, and ecs.sgn7i-vws-m8s.4xlarge. The following table describes the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance type	vCPUs	Memory (GiB)	GPUs	GPU memory	Network baseline/burst bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.sgn7i-vws-m2.xlarge	4	15.5	NVIDIA A10 * 1/12	24GB * 1/12	1.5/5	500,000	4	2	2	1

ecs.sgn7i-vws-m4.2xlarge	8	31	NVIDIA A10 * 1/6	24GB * 1/6	2.5/10	1,000,000	4	4	6	1
ecs.sgn7i-vws-m8.4xlarge	16	62	NVIDIA A10 * 1/3	24GB * 1/3	5/20	2,000,000	8	4	10	1
ecs.sgn7i-vws-m2s.xlarge	4	8	NVIDIA A10 * 1/12	24GB * 1/12	1.5/5	500,000	4	2	2	1
ecs.sgn7i-vws-m4s.2xlarge	8	16	NVIDIA A10 * 1/6	24GB * 1/6	2.5/10	1,000,000	4	4	6	1
ecs.sgn7i-vws-m8s.4xlarge	16	32	NVIDIA A10 * 1/3	24GB * 1/3	5/20	2,000,000	8	4	10	1

Note

The **GPU** column in the preceding table indicates the GPU model and GPU slicing information for each instance type. Each GPU can be sliced into multiple GPU partitions, and each GPU partition can be allocated as a vGPU to an instance. Example:

`NVIDIA A10 * 1/12`. `NVIDIA A10` is the GPU model. `1/12` indicates that a GPU is sliced into 12 GPU partitions, and each GPU partition can be allocated as a vGPU to an instance.

vgn7i-vws, vGPU-accelerated instance family

- **Introduction**:
  - This instance family uses the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude. This way, data storage and model loading can be performed more quickly.
  - This instance family comes with an NVIDIA GRID vWS license and provides certified graphics acceleration capabilities for CAD software to meet the requirements of professional graphic design. Instances of this instance family can serve as lightweight GPU-accelerated compute-optimized instances to reduce the costs of small-scale AI inference tasks.
- **Supported scenarios**:
  - Concurrent AI inference tasks that require high-performance CPUs, memory, and GPUs, such as image recognition, speech recognition, and behavior identification
  - Compute-intensive graphics processing tasks that require high-performance 3D graphics virtualization capabilities, such as remote graphic design and cloud gaming
  - 3D modeling in fields that require the use of Ice Lake processors, such as animation and film production, cloud gaming, and mechanical design
- **Compute**:

- Uses NVIDIA A10 GPUs that have the following features:
  - Innovative NVIDIA Ampere architecture
  - Support for acceleration features, such as vGPU, RTX, and TensorRT, to provide diversified business support
- Uses 2.9 GHz Intel® Xeon® Scalable (Ice Lake) processors that deliver an all-core turbo frequency of 3.5 GHz.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

The vgn7i-vws instance family includes the following instance types: ecs.vgn7i-vws-m4.xlarge, ecs.vgn7i-vws-m8.2xlarge, ecs.vgn7i-vws-m12.3xlarge, and ecs.vgn7i-vws-m24.7xlarge. Expand the following section to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	GPUs	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.vgn7i-vws-m4.xlarge	4	30	NVIDIA A10 * 1/6	24GB * 1/6	3	1,000,000	4	4	10	1
ecs.vgn7i-vws-m8.2xlarge	10	62	NVIDIA A10 * 1/3	24GB * 1/3	5	2,000,000	8	6	10	1
ecs.vgn7i-vws-m12.3xlarge	14	93	NVIDIA A10 * 1/2	24GB * 1/2	8	3,000,000	8	6	15	1
ecs.vgn7i-vws-m24.7xlarge	30	186	NVIDIA A10 * 1	24GB * 1	16	6,000,000	12	8	30	1

> **Note**
>
> The **GPU** column in the preceding table indicates the GPU model and GPU slicing information for each instance type. Each GPU can be sliced into multiple GPU partitions, and each GPU partition can be allocated as a vGPU to an instance. Example:
>
> `NVIDIA A10 * 1/6`. `NVIDIA A10` is the GPU model. `1/6` indicates that a GPU is sliced into six GPU partitions, and each GPU partition can be allocated as a vGPU to an instance.

vgn6i-vws, vGPU-accelerated instance family

> **Important**
>
> - In light of the NVIDIA GRID driver upgrade, Alibaba Cloud upgrades the vgn6i instance family to the vgn6i-vws instance family. The vgn6i-vws instance family uses the latest NVIDIA GRID driver and provides an NVIDIA GRID vWS license. To apply for free images for which the NVIDIA GRID driver is pre-installed, submit a ticket.
> - To use other public images or custom images that do not contain an NVIDIA GRID driver, submit a ticket to apply for the GRID driver file and install the NVIDIA GRID driver. Alibaba Cloud does not charge additional license fees for the GRID driver.

- **Supported scenarios**:
  - Real-time rendering for cloud gaming
  - Real-time rendering for Augmented Reality (AR) and Virtual Reality (VR) applications
  - AI (deep learning and machine learning) inference for elastic Internet service deployment
  - Educational environment of deep learning
  - Modeling experiment environment of deep learning
- **Compute**:
  - Uses NVIDIA T4 GPUs.
  - Uses vGPUs.
    - Supports the 1/4 and 1/2 compute capacity of NVIDIA Tesla T4 GPUs.
    - Supports 4 GB and 8 GB of GPU memory.
  - Offers a CPU-to-memory ratio of 1:5.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

The vgn6i-vws instance family includes the following instance types: ecs.vgn6i-m4-vws.xlarge, ecs.vgn6i-m8-vws.2xlarge, and ecs.vgn6i-m16-vws.5xlarge. Expand the following section to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	GPUs	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.vgn6i-m4-vws.xlarge	4	23	NVIDIA T4 * 1/4	16GB * 1/4	2	500,000	4/2	3	10	1
ecs.vgn6i-m8-vws.2xlarge	10	46	NVIDIA T4 * 1/2	16GB * 1/2	4	800,000	8/2	4	10	1
ecs.vgn6i-m16-vws.5xlarge	20	92	NVIDIA T4 * 1	16GB * 1	7.5	1,200,000	6	4	10	1

Note

The **GPU** column in the preceding table indicates the GPU model and GPU slicing information for each instance type. Each GPU can be sliced into multiple GPU partitions, and each GPU partition can be allocated as a vGPU to an instance. Example:

`NVIDIA T4 * 1/4`. `NVIDIA T4` is the GPU model. `1/4` indicates that a GPU is sliced into four GPU partitions, and each GPU partition can be allocated as a vGPU to an instance.

gn8v, GPU-accelerated compute-optimized instance family

This instance family is available only in specific regions, including regions outside China. To use the instance family, contact Alibaba Cloud sales personnel.

- **Introduction**: This instance family is an 8th-generation GPU-accelerated compute-optimized instance family provided by Alibaba Cloud for AI model training and the inference tasks of ultra-large models. This instance family consists of multiple instance types that provide one, two, four, or eight GPUs per instance.
- **Supported scenarios**:
  - Multi-GPU parallel inference computing for large language models (LLMs) that have more than 70 billion parameters
  - Traditional AI model training and autonomous driving training, for which each GPU delivers computing power of up to 39.5 TFLOPS in the single-precision floating-point format (FP32)
  - Small and medium-sized model training scenarios that leverage the NVLink connections among the eight GPUs
- **Benefits and positioning**:

- **High-speed and large-capacity GPU memory**: Each GPU is equipped with 96 GB of HBM3E memory and delivers up to 4 TB/s of memory bandwidth, which greatly accelerates model training and inference.
- **High bandwidth between GPUs**: Multiple GPUs are interconnected by using 900 GB/s NVLink connections. The efficiency of multi-GPU training and inference is much higher than that of previous generations of GPU-accelerated instances.
- **Quantization of large models**: This instance family supports computing power in the 8-bit floating point format (FP8) and optimizes computing power for large-scale parameter training and inference. This significantly improves the computing speed of training and inference and reduces memory usage.
- **High security**: This instance family supports confidential computing capabilities that cover the full link of model inference tasks. The capabilities include CPU-based Intel Trust Domain Extensions (TDX) confidential computing and GPU-based NVIDIA Confidential Computing (CC). The confidential computing capabilities ensure the security of user inference data and enterprise models in model inference and training.

- **Compute**:
  - Uses the latest Cloud Infrastructure Processing Unit (CIPU) 1.0 processors.
    - Decouples computing capabilities from storage capabilities, allowing you to flexibly select storage resources based on your business requirements.
    - Provides bare metal capabilities to support peer-to-peer (P2P) communication between GPU-accelerated instances.
  - Uses the 4th-generation Intel Xeon Scalable processors that deliver a base frequency of up to 2.8 GHz and an all-core turbo frequency of up to 3.1 GHz.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, and elastic ephemeral disks (EEDs).
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports the Jumbo Frames feature. For more information, see Jumbo Frames.
  - Provides ultra-high network performance with a packet forwarding rate of up to 30,000,000 pps (for instances equipped with eight GPUs).
  - Supports elastic RDMA interfaces (ERIs).

    > **Note**
    >
    > For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.

The gn8v instance family includes the following instance types: ecs.gn8v.4xlarge, ecs.gn8v.6xlarge, ecs.gn8v-2x.8xlarge, ecs.gn8v-4x.8xlarge, ecs.gn8v-2x.12xlarge, ecs.gn8v-8x.16xlarge, ecs.gn8v-4x.24xlarge, and ecs.gn8v-8x.48xlarge. Expand the following section to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	GPU memory	Network baseline bandwidth (Gbit/s)	ENIs	NIC queues per primary ENI	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Maximum disks	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)

ecs.gn8v.4xlarge	16	96	96GB * 1	12	8	16	30	30	17	100,000	0.75
ecs.gn8v.6xlarge	24	128	96GB * 1	15	8	24	30	30	17	120,000	0.937
ecs.gn8v-2x.8xlarge	32	192	96GB * 2	20	8	32	30	30	25	200,000	1.25
ecs.gn8v-4x.8xlarge	32	384	96GB * 4	20	8	32	30	30	25	200,000	1.25
ecs.gn8v-2x.12xlarge	48	256	96GB * 2	25	8	48	30	30	33	300,000	1.50
ecs.gn8v-8x.16xlarge	64	768	96GB * 8	32	8	64	30	30	33	360,000	2.5
ecs.gn8v-4x.24xlarge	96	512	96GB * 4	50	15	64	30	30	49	500,000	3
ecs.gn8v-8x.48xlarge	192	1,024	96GB * 8	100	15	64	50	50	65	1,000,000	6

gn8is, GPU-accelerated compute-optimized instance family

This instance family is available only in specific regions, including regions outside China. To use the instance family, contact Alibaba Cloud sales personnel.

- **Introduction**: This instance family is an 8th-generation GPU-accelerated compute-optimized instance family provided by Alibaba Cloud in response to the recent developments in the AI generation field. This instance family consists of multiple instance types that provide one, two, four, or eight GPUs per instance and have different CPU-to-GPU ratios to fit various use cases.
- **Benefits and positioning**:
  - **Graphic processing**: This instance family uses high-frequency 5th-generation Intel Xeon Scalable processors to provide sufficient CPU capacity for smooth graphics rendering and design in 3D modeling scenarios.

- **Inference tasks**: This instance family uses innovative GPUs, each with 48 GB of memory, which accelerate inference tasks and support the FP8 floating-point format. You can use this instance family together with Container Service for Kubernetes (ACK) to support the inference of various AI-generated content (AIGC) models and especially accommodate inference tasks for LLMs that have less than 70 billion parameters.

- **Supported scenarios**:
  - Animation, special effects for film and television, and rendering
  - Generation of AIGC images and inference of LLMs
  - Other general-purpose AI recognition, image recognition, and speech recognition scenarios
- **Compute**:
  - Uses innovative GPUs that have the following features:
    - Support for acceleration features, such as TensorRT, and the FP8 floating-point format to improve LLM inference performance.
    - Up to 48 GB of memory per GPU and support for the inference of 70B or larger LLMs on a single instance with multiple GPUs.
    - Improved graphic processing capabilities. For example, after you install a GRID driver on a gn8is instance by using Cloud Assistant or an Alibaba Cloud Marketplace image, the instance can provide graphic processing performance twice that of a 7th-generation instance.
  - Uses the latest high-frequency Intel® Xeon® processors that deliver an all-core turbo frequency of 3.9 GHz to meet complex 3D modeling requirements.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, and EEDs.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Supports ERIs.

    > **Note**
    >
    > For information about how to use ERIs, see Configure eRDMA on an enterprise-level instance.

The gn8is instance family includes the following instance types: ecs.gn8is.2xlarge, ecs.gn8is.4xlarge, ecs.gn8is-2x.8xlarge, ecs.gn8is-4x.16xlarge, and ecs.gn8is-8x.32xlarge. Expand the following section to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	GPU memory	Network baseline bandwidth (Gbit/s)	ENIs	NIC queues per primary ENI	Private IPv4 addresses per ENI	IPv6 addresses per ENI	Maximum disks	Disk baseline IOPS	Disk baseline bandwidth (Gbit/s)
ecs.gn8is.2xlarge	8	64	48GB * 1	8	4	8	15	15	17	60,000	0.75
ecs.gn8is.4xlarge	16	128	48GB * 1	16	8	16	30	30	17	120,000	1.25

ecs.gn8is-2x.8xlarge	32	256	48GB * 2	32	8	32	30	30	33	250,000	2
ecs.gn8is-4x.16xlarge	64	512	48GB * 4	64	8	64	30	30	33	450,000	4
ecs.gn8is-8x.32xlarge	128	1,024	48GB * 8	100	15	64	50	50	65	900,000	8

gn7e, GPU-accelerated compute-optimized instance family

Features:

- **Introduction**:
  - This instance family allows you to select instance types that provide different numbers of GPUs and CPUs to meet your business requirements in AI use cases.
  - This instance family uses the third-generation SHENLONG architecture and doubles the average bandwidths of virtual private clouds (VPCs), networks, and disks compared with instance families of the previous generation.
- **Supported scenarios**:
  - Small- and medium-scale AI training
  - High-performance computing (HPC) business accelerated by using Compute Unified Device Architecture (CUDA)
  - AI inference tasks that require high GPU processing capabilities or large amounts of GPU memory
  - Deep learning applications, such as training applications of AI algorithms used in image classification, autonomous vehicles, and speech recognition
  - Scientific computing applications that require robust GPU computing capabilities, such as computational fluid dynamics, computational finance, molecular dynamics, and environmental analytics

> **Important**
>
> When you use AI training services that feature a high communication load, such as transformer models, you must enable NVLink for GPU-to-GPU communication. Otherwise, data may be damaged due to unpredictable failures that are caused by large-scale data transmission over Peripheral Component Interconnect Express (PCIe) links. If you do not understand the topology of the communication links that are used for AI training services, submit a ticket to obtain technical support.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

The gn7e instance family includes the following instance types: ecs.gn7e-c16g1.4xlarge, ecs.gn7e-c16g1.8xlarge, ecs.gn7e-c16g1.16xlarge, and ecs.gn7e-c16g1.32xlarge. Expand the following section to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.gn7e-c16g1.4xlarge	16	125	80GB * 1	8	3,000,000	8	8	10	1
ecs.gn7e-c16g1.8xlarge	32	250	80GB * 2	16	6,000,000	16	8	10	1
ecs.gn7e-c16g1.16xlarge	64	500	80GB * 4	32	12,000,000	32	8	10	1
ecs.gn7e-c16g1.32xlarge	128	1,000	80GB * 8	64	24,000,000	32	16	15	1

gn7i, GPU-accelerated compute-optimized instance family

- **Introduction**: This instance family uses the third-generation SHENLONG architecture to provide predictable and consistent ultra-high performance. This instance family utilizes fast path acceleration on chips to improve storage performance, network performance, and computing stability by an order of magnitude.
- **Supported scenarios**:
  - Concurrent AI inference tasks that require high-performance CPUs, memory, and GPUs, such as image recognition, speech recognition, and behavior identification
  - Compute-intensive graphics processing tasks that require high-performance 3D graphics virtualization capabilities, such as remote graphic design and cloud gaming
- **Compute**:
  - Uses NVIDIA A10 GPUs that have the following features:
    - Innovative NVIDIA Ampere architecture
    - Support for acceleration features, such as RTX and TensorRT
  - Uses 2.9 GHz Intel® Xeon® Scalable (Ice Lake) processors that deliver an all-core turbo frequency of 3.5 GHz.
  - Provides up to 752 GiB of memory, which is much larger than the memory sizes of the gn6i instance family.
- **Storage**:

- Is an instance family in which all instances are I/O optimized.
- Supports ESSDs and ESSD AutoPL disks.

- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

The gn7i instance family includes the following instance types: ecs.gn7i-c8g1.2xlarge, ecs.gn7i-c16g1.4xlarge, ecs.gn7i-c32g1.8xlarge, ecs.gn7i-c32g1.16xlarge, ecs.gn7i-c32g1.32xlarge, ecs.gn7i-c48g1.12xlarge, ecs.gn7i-c56g1.14xlarge, ecs.gn7i-2x.8xlarge, ecs.gn7i-4x.8xlarge, ecs.gn7i-4x.16xlarge, ecs.gn7i-8x.32xlarge, and ecs.gn7i-8x.16xlarge. Expand the following section to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	GPUs	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.gn7i-c8g1.2xlarge	8	30	NVIDIA A10 * 1	24GB * 1	16	1,600,000	8	4	15	15
ecs.gn7i-c16g1.4xlarge	16	60	NVIDIA A10 * 1	24GB * 1	16	3,000,000	8	8	30	30
ecs.gn7i-c32g1.8xlarge	32	188	NVIDIA A10 * 1	24GB * 1	16	6,000,000	12	8	30	30
ecs.gn7i-c32g1.16xlarge	64	376	NVIDIA A10 * 2	24GB * 2	32	12,000,000	16	15	30	30
ecs.gn7i-c32g1.32xlarge	128	752	NVIDIA A10 * 4	24GB * 4	64	24,000,000	32	15	30	30
ecs.gn7i-c48g1.12xlarge	48	310	NVIDIA A10 * 1	24GB * 1	16	9,000,000	16	8	30	30

ecs.gn7i-c56g1.14xlarge	56	346	NVIDIA A10 * 1	24GB * 1	16	12,000,000	16	12	30	30
ecs.gn7i-2x.8xlarge	32	128	NVIDIA A10 * 2	24GB * 2	16	6,000,000	16	8	30	30
ecs.gn7i-4x.8xlarge	32	128	NVIDIA A10 * 4	24GB * 4	16	6,000,000	16	8	30	30
ecs.gn7i-4x.16xlarge	64	256	NVIDIA A10 * 4	24GB * 4	32	12,000,000	32	8	30	30
ecs.gn7i-8x.32xlarge	128	512	NVIDIA A10 * 8	24GB * 8	64	24,000,000	32	16	30	30
ecs.gn7i-8x.16xlarge	64	256	NVIDIA A10 * 8	24GB * 8	32	12,000,000	32	8	30	30

Important

You can change the following instance types only to ecs.gn7i-c8g1.2xlarge or ecs.gn7i-c16g1.4xlarge: ecs.gn7i-2x.8xlarge, ecs.gn7i-4x.8xlarge, ecs.gn7i-4x.16xlarge, ecs.gn7i-8x.32xlarge, and ecs.gn7i-8x.16xlarge.

gn7s, GPU-accelerated compute-optimized instance family

To use the gn7s instance family, submit a ticket.

- **Introduction**:
  - This instance family uses the latest Intel Ice Lake processors and NVIDIA A30 GPUs that are based on NVIDIA Ampere architecture. You can select instance types that comprise appropriate mixes of GPUs and vCPUs to meet your business requirements in AI scenarios.
  - This instance family uses the third-generation SHENLONG architecture and doubles the average bandwidths of VPCs, networks, and disks compared with instance families of the previous generation.
- **Supported scenarios**: concurrent AI inference tasks that require high-performance CPUs, memory, and GPUs, such as image recognition, speech recognition, and behavior identification.
- **Compute**:

- Uses NVIDIA A30 GPUs that have the following features:
  - Innovative NVIDIA Ampere architecture
  - Support for the multi-instance GPU (MIG) feature and acceleration features (based on second-generation Tensor cores) to provide diversified business support
- Uses 2.9 GHz Intel® Xeon® Scalable (Ice Lake) processors that deliver an all-core turbo frequency of 3.5 GHz.
- Improves memory sizes significantly from instance families of the previous generation.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

The gn7s instance family includes the following instance types: ecs.gn7s-c8g1.2xlarge, ecs.gn7s-c16g1.4xlarge, ecs.gn7s-c32g1.8xlarge, ecs.gn7s-c32g1.16xlarge, ecs.gn7s-c32g1.32xlarge, ecs.gn7s-c48g1.12xlarge, and ecs.gn7s-c56g1.14xlarge. Expand the following section to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	GPUs	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Private IPv4 addresses per ENI	IPv6 addresses per ENI	NIC queues	ENIs
ecs.gn7s-c8g1.2xlarge	8	60	NVIDIA A30 * 1	24GB * 1	16	6,000,000	5	1	12	8
ecs.gn7s-c16g1.4xlarge	16	120	NVIDIA A30 * 1	24GB * 1	16	6,000,000	5	1	12	8
ecs.gn7s-c32g1.8xlarge	32	250	NVIDIA A30 * 1	24GB * 1	16	6,000,000	5	1	12	8
ecs.gn7s-c32g1.16xlarge	64	500	NVIDIA A30 * 2	24GB * 2	32	12,000,000	5	1	16	15
ecs.gn7s-c32g1.32xlarge	128	1,000	NVIDIA A30 * 4	24GB * 4	64	24,000,000	10	1	32	15

ecs.gn7s-c48g1.12xlarge	48	380	NVIDIA A30 * 1	24GB * 1	16	6,000,000	8	1	12	8
ecs.gn7s-c56g1.14xlarge	56	440	NVIDIA A30 * 1	24GB * 1	16	6,000,000	8	1	12	8

gn7, GPU-accelerated compute-optimized instance family

- **Supported scenarios**:
  - Deep learning applications, such as training applications of AI algorithms used in image classification, autonomous vehicles, and speech recognition
  - Scientific computing applications that require robust GPU computing capabilities, such as computational fluid dynamics, computational finance, molecular dynamics, and environmental analytics
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs and ESSD AutoPL disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

The gn7 instance family includes the following instance types: ecs.gn7-c12g1.3xlarge, ecs.gn7-c13g1.13xlarge, and ecs.gn7-c13g1.26xlarge. Expand the following section to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.gn7-c12g1.3xlarge	12	94	40GB * 1	4	2,500,000	4	8	10	1
ecs.gn7-c13g1.13xlarge	52	378	40GB * 4	16	9,000,000	16	8	30	30
ecs.gn7-c13g1.26xlarge	104	756	40GB * 8	30	18,000,000	16	15	10	1

gn6i, GPU-accelerated compute-optimized instance family

- **Supported scenarios**:
  - AI (deep learning and machine learning) inference for computer vision, speech recognition, speech synthesis, natural language processing (NLP), machine translation, and recommendation systems
  - Real-time rendering for cloud gaming
  - Real-time rendering for AR and VR applications
  - Graphics workstations or graphics-heavy computing
  - GPU-accelerated databases
  - High-performance computing
- **Compute**:
  - Uses NVIDIA T4 GPUs that have the following features:
    - Innovative NVIDIA Turing architecture
    - 16 GB of memory (320 GB/s bandwidth) per GPU
    - 2,560 CUDA cores per GPU
    - Up to 320 Turing Tensor cores per GPU
    - Mixed-precision Tensor cores that support 65 FP16 TFLOPS, 130 INT8 TOPS, and 260 INT4 TOPS
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors.
- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

The gn6i instance family includes the following instance types: ecs.gn6i-c4g1.xlarge, ecs.gn6i-c8g1.2xlarge, ecs.gn6i-c16g1.4xlarge, ecs.gn6i-c24g1.6xlarge, ecs.gn6i-c40g1.10xlarge, ecs.gn6i-c24g1.12xlarge, and ecs.gn6i-c24g1.24xlarge. Expand the following section to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	GPUs	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Disk baseline IOPS	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.gn6i-c4g1.xlarge	4	15	NVIDIA T4 * 1	16GB * 1	4	500,000	None	2	2	10	1

ecs.gn6i-c8g1.2xlarge	8	31	NVIDIA T4 * 1	16GB * 1	5	800,000	None	2	2	10	1
ecs.gn6i-c16g1.4xlarge	16	62	NVIDIA T4 * 1	16GB * 1	6	1,000,000	None	4	3	10	1
ecs.gn6i-c24g1.6xlarge	24	93	NVIDIA T4 * 1	16GB * 1	7.5	1,200,000	None	6	4	10	1
ecs.gn6i-c40g1.10xlarge	40	155	NVIDIA T4 * 1	16GB * 1	10	1,600,000	None	16	10	10	1
ecs.gn6i-c24g1.12xlarge	48	186	NVIDIA T4 * 2	16GB * 2	15	2,400,000	None	12	6	10	1
ecs.gn6i-c24g1.24xlarge	96	372	NVIDIA T4 * 4	16GB * 4	30	4,800,000	250,000	24	8	10	1

gn6e, GPU-accelerated compute-optimized instance family

- **Supported scenarios**:
  - Deep learning applications, such as training and inference applications of AI algorithms used in image classification, autonomous vehicles, and speech recognition
  - Scientific computing applications, such as computational fluid dynamics, computational finance, molecular dynamics, and environmental analytics
- **Compute**:
  - Uses NVIDIA V100 GPUs that each have 32 GB of GPU memory and support NVLink.
  - Uses NVIDIA V100 GPUs (SXM2-based) that have the following features:
    - Innovative NVIDIA Volta architecture
    - 32 GB of HBM2 memory (900 GB/s bandwidth) per GPU
    - 5,120 CUDA cores per GPU
    - 640 Tensor cores per GPU
    - Up to six NVLink bidirectional connections per GPU, each of which provides a bandwidth of 25 Gbit/s in each direction for a total bandwidth of 300 Gbit/s (6 × 25 × 2 = 300)
  - Offers a CPU-to-memory ratio of 1:8.

- Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

The gn6e instance family includes the following instance types: ecs.gn6e-c12g1.3xlarge, ecs.gn6e-c12g1.6xlarge, ecs.gn6e-c12g1.12xlarge, and ecs.gn6e-c12g1.24xlarge. Expand the following section to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	GPUs	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.gn6e-c12g1.3xlarge	12	92	NVIDIA V100 * 1	32GB * 1	5	800,000	8	6	10	1
ecs.gn6e-c12g1.6xlarge	24	182	NVIDIA V100 * 2	32GB * 2	8	1,200,000	8	8	20	1
ecs.gn6e-c12g1.12xlarge	48	368	NVIDIA V100 * 4	32GB * 4	16	2,400,000	8	8	20	1
ecs.gn6e-c12g1.24xlarge	96	736	NVIDIA V100 * 8	32GB * 8	32	4,800,000	16	8	20	1

gn6v, GPU-accelerated compute-optimized instance family

- **Supported scenarios**:
  - Deep learning applications, such as training and inference applications of AI algorithms used in image classification, autonomous vehicles, and speech recognition
  - Scientific computing applications, such as computational fluid dynamics, computational finance, molecular dynamics, and environmental analytics
- **Compute**:
  - Uses NVIDIA V100 GPUs.

- Uses NVIDIA V100 GPUs (SXM2-based) that have the following features:
  - Innovative NVIDIA Volta architecture
  - 16 GB of HBM2 memory (900 GB/s bandwidth) per GPU
  - 5,120 CUDA cores per GPU
  - 640 Tensor cores per GPU
  - Up to six NVLink bidirectional connections per GPU, each of which provides a bandwidth of 25 Gbit/s in each direction for a total bandwidth of 300 Gbit/s (6 × 25 × 2 = 300)
- Offers a CPU-to-memory ratio of 1:4.
- Uses 2.5 GHz Intel® Xeon® Platinum 8163 (Skylake) processors.

- **Storage**:
  - Is an instance family in which all instances are I/O optimized.
  - Supports ESSDs, ESSD AutoPL disks, standard SSDs, and ultra disks.
- **Network**:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

The gn6v instance family includes the following instance types: ecs.gn6v-c8g1.2xlarge, ecs.gn6v-c8g1.4xlarge, ecs.gn6v-c8g1.8xlarge, ecs.gn6v-c8g1.16xlarge, and ecs.gn6v-c10g1.20xlarge. Expand the following section to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	GPUs	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	Disk baseline IOPS	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.gn6v-c8g1.2xlarge	8	32	NVIDIA V100 * 1	16GB * 1	2.5	800,000	None	4	4	10	1
ecs.gn6v-c8g1.4xlarge	16	64	NVIDIA V100 * 2	16GB * 2	5	1,000,000	None	4	8	20	1
ecs.gn6v-c8g1.8xlarge	32	128	NVIDIA V100 * 4	16GB * 4	10	2,000,000	None	8	8	20	1
ecs.gn6v-c8g1.16xlarge	64	256	NVIDIA V100 * 8	16GB * 8	20	2,500,000	None	16	8	20	1

ecs.gn6v-c10g1.20xlarge	82	336	NVIDIA V100 * 8	16GB * 8	32	4,500,000	250,000	16	8	20	1

gn5, GPU-accelerated compute-optimized instance family

- **Supported scenarios**:
  - Deep learning
  - Scientific computing applications, such as computational fluid dynamics, computational finance, genomics, and environmental analytics
  - Server-side GPU compute workloads, such as high-performance computing, rendering, and multi-media encoding and decoding
- **Compute**:
  - Uses NVIDIA P100 GPUs.
  - Offers multiple CPU-to-memory ratios.
  - Uses 2.5 GHz Intel® Xeon® E5-2682 v4 (Broadwell) processors.
- **Storage**:
  - Supports high-performance local Non-Volatile Memory Express (NVMe) SSDs.
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks.
- **Network**:
  - Supports only IPv4.
  - Provides high network performance based on large computing capacity.

The gn5 instance family includes the following instance types: ecs.gn5-c4g1.xlarge, ecs.gn5-c8g1.2xlarge, ecs.gn5-c4g1.2xlarge, ecs.gn5-c8g1.4xlarge, ecs.gn5-c28g1.7xlarge, ecs.gn5-c8g1.8xlarge, ecs.gn5-c28g1.14xlarge, and ecs.gn5-c8g1.14xlarge. Expand the following section to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	Local storage (GiB)	GPUs	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI
ecs.gn5-c4g1.xlarge	4	30	440	NVIDIA P100 * 1	16GB * 1	3	300,000	1	3	10
ecs.gn5-c8g1.2xlarge	8	60	440	NVIDIA P100 * 1	16GB * 1	3	400,000	1	4	10

ecs.gn5-c4g1.2xlarge	8	60	880	NVIDIA P100 * 2	16GB * 2	5	1,000,000	2	4	10
ecs.gn5-c8g1.4xlarge	16	120	880	NVIDIA P100 * 2	16GB * 2	5	1,000,000	4	8	20
ecs.gn5-c28g1.7xlarge	28	112	440	NVIDIA P100 * 1	16GB * 1	5	1,000,000	8	8	20
ecs.gn5-c8g1.8xlarge	32	240	1,760	NVIDIA P100 * 4	16GB * 4	10	2,000,000	8	8	20
ecs.gn5-c28g1.14xlarge	56	224	880	NVIDIA P100 * 2	16GB * 2	10	2,000,000	14	8	20
ecs.gn5-c8g1.14xlarge	54	480	3,520	NVIDIA P100 * 8	16GB * 8	25	4,000,000	14	8	20

gn5i, GPU-accelerated compute-optimized instance family

- **Supported scenarios**: server-side GPU compute workloads, such as deep learning inference and multi-media encoding and decoding.
- Compute:
  - Uses NVIDIA P4 GPUs.
  - Offers a CPU-to-memory ratio of 1:4.
  - Uses 2.5 GHz Intel® Xeon® E5-2682 v4 (Broadwell) processors.
- Storage:
  - Is an instance family in which all instances are I/O optimized.
  - Supports standard SSDs and ultra disks.
- Network:
  - Supports IPv4 and IPv6. For information about IPv6 communication, see IPv6 communication.
  - Provides high network performance based on large computing capacity.

The gn5i instance family includes the following instance types: ecs.gn5i-c2g1.large, ecs.gn5i-c4g1.xlarge, ecs.gn5i-c8g1.2xlarge, ecs.gn5i-c16g1.4xlarge, ecs.gn5i-c16g1.8xlarge, and ecs.gn5i-c28g1.14xlarge. Expand the following section to see a table describing the specifications of each instance type in this instance family. For information about the metrics of instance types, see Instance type metrics.

Instance types

Instance type	vCPUs	Memory (GiB)	GPUs	GPU memory	Network baseline bandwidth (Gbit/s)	Packet forwarding rate (pps)	NIC queues	ENIs	Private IPv4 addresses per ENI	IPv6 addresses per ENI
ecs.gn5i-c2g1.large	2	8	NVIDIA P4 * 1	8GB * 1	1	100,000	2	2	6	1
ecs.gn5i-c4g1.xlarge	4	16	NVIDIA P4 * 1	8GB * 1	1.5	200,000	2	3	10	1
ecs.gn5i-c8g1.2xlarge	8	32	NVIDIA P4 * 1	8GB * 1	2	400,000	4	4	10	1
ecs.gn5i-c16g1.4xlarge	16	64	NVIDIA P4 * 1	8GB * 1	3	800,000	4	8	20	1
ecs.gn5i-c16g1.8xlarge	32	128	NVIDIA P4 * 2	8GB * 2	6	1,200,000	8	8	20	1
ecs.gn5i-c28g1.14xlarge	56	224	NVIDIA P4 * 2	8GB * 2	10	2,000,000	14	8	20	1