

Apache CloudStack 4.1.0

Guia do Administrador do CloudStack



Apache CloudStack

Apache CloudStack 4.1.0 Guia do Administrador do CloudStack

Autor

Apache CloudStack

Licensed to the Apache Software Foundation (ASF) under one or more contributor license agreements. See the NOTICE file distributed with this work for additional information regarding copyright ownership. The ASF licenses this file to you under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Apache CloudStack is an effort undergoing incubation at The Apache Software Foundation (ASF).

Incubation is required of all newly accepted projects until a further review indicates that the infrastructure, communications, and decision making process have stabilized in a manner consistent with other successful ASF projects. While incubation status is not necessarily a reflection of the completeness or stability of the code, it does indicate that the project has yet to be fully endorsed by the ASF.

Guia de administração do CloudStack.

1. Conceitos	1
1.1. O que é o CloudStack?	1
1.2. O que o CloudStack pode fazer?	1
1.3. Visão geral da arquitetura de implementação	2
1.3.1. Visão geral da arquitetura de implementação	3
1.3.2. Visão geral da infraestrutura de nuvem	4
1.3.3. Visão geral de serviços de rede	5
2. Conceitos de infraestrutura de nuvem	7
2.1. About Regions	7
2.2. Sobre zonas	7
2.3. Sobre pods	9
2.4. Sobre clusters	9
2.5. Sobre hosts	10
2.6. Sobre storage primária	11
2.7. Sobre storage secundária	11
2.8. Sobre redes físicas	12
2.8.1. Tipos de tráfego de rede da zona básica	12
2.8.2. Endereços IP de hóspedes na zona básica	13
2.8.3. Tipos de tráfego de rede da zona avançada	13
2.8.4. Endereços IP de hóspedes na zona avançada	14
2.8.5. Endereços IP públicos na zona avançada	14
2.8.6. Endereços IP reservados pelo sistema	14
3. Accounts	17
3.1. Accounts, Users, and Domains	17
3.2. Using an LDAP Server for User Authentication	17
3.2.1. Example LDAP Configuration Commands	18
3.2.2. Search Base	18
3.2.3. Query Filter	19
3.2.4. Search User Bind DN	19
3.2.5. SSL Keystore Path and Password	20
4. User Services Overview	21
4.1. Service Offerings, Disk Offerings, Network Offerings, and Templates	21
5. Interface do usuário	23
5.1. Login na interface de usuário	23
5.1.1. End User's UI Overview	23
5.1.2. Root Administrator's UI Overview	23
5.1.3. Fazendo login como o administrador root	24
5.1.4. Changing the Root Password	25
5.2. Usando as chaves SSH para autenticação.	25
5.2.1. Criando um template de instância que suporta chaves SSH	25
5.2.2. Criando o par de chaves SSH	26
5.2.3. Criando uma instância	27
5.2.4. Fazendo login usando o par de chaves SSH	27
5.2.5. Resetting SSH Keys	27
6. Using Projects to Organize Users and Resources	29
6.1. Overview of Projects	29
6.2. Configuring Projects	29
6.2.1. Setting Up Invitations	29
6.2.2. Setting Resource Limits for Projects	30
6.2.3. Setting Project Creator Permissions	31
6.3. Creating a New Project	32

6.4. Adding Members to a Project	32
6.4.1. Sending Project Membership Invitations	32
6.4.2. Adding Project Members From the UI	33
6.5. Accepting a Membership Invitation	33
6.6. Suspending or Deleting a Project	34
6.7. Using the Project View	34
7. Passos para provisionamento de sua infraestrutura de nuvem	35
7.1. Visão geral dos passos de provisionamento	35
7.2. Adding Regions (optional)	36
7.2.1. The First Region: The Default Region	36
7.2.2. Adding a Region	36
7.2.3. Adding Third and Subsequent Regions	37
7.2.4. Deleting a Region	39
7.3. Adicionando uma zona	39
7.3.1. Configuração de zona básica	40
7.3.2. Advanced Zone Configuration	44
7.4. Adicionando um pod	48
7.5. Adicionando um cluster	49
7.5.1. Add Cluster: KVM or XenServer	49
7.5.2. Add Cluster: vSphere	49
7.6. Adding a Host	51
7.6.1. Adding a Host (XenServer or KVM)	52
7.6.2. Adding a Host (vSphere)	53
7.7. Adicionar Storage Primário	54
7.7.1. System Requirements for Primary Storage	54
7.7.2. Adding Primary Storage	54
7.8. Adicionar Storage Secundário	55
7.8.1. System Requirements for Secondary Storage	55
7.8.2. Adding Secondary Storage	55
7.9. Initialize and Test	56
8. Ofertas de serviços	59
8.1. Compute and Disk Service Offerings	59
8.1.1. Creating a New Compute Offering	59
8.1.2. Creating a New Disk Offering	60
8.1.3. Modifying or Deleting a Service Offering	61
8.2. System Service Offerings	61
8.2.1. Creating a New System Service Offering	61
8.3. Network Throttling	62
8.4. Changing the Default System Offering for System VMs	64
9. Setting Up Networking for Users	65
9.1. Overview of Setting Up Networking for Users	65
9.2. About Virtual Networks	65
9.2.1. Isolated Networks	65
9.2.2. Shared Networks	65
9.2.3. Runtime Allocation of Virtual Network Resources	66
9.3. Provedores de Serviços de Rede	66
9.4. Oferta de Rede	67
9.4.1. Creating a New Network Offering	68
10. Working With Virtual Machines	71
10.1. About Working with Virtual Machines	71
10.2. Best Practices for Virtual Machines	71
10.3. VM Lifecycle	72

10.4. Creating VMs	72
10.5. Accessing VMs	74
10.6. Stopping and Starting VMs	74
10.7. Changing the VM Name, OS, or Group	75
10.8. Changing the Service Offering for a VM	75
10.9. Moving VMs Between Hosts (Manual Live Migration)	75
10.10. Deleting VMs	76
10.11. Working with ISOs	76
10.11.1. Adding an ISO	77
10.11.2. Attaching an ISO to a VM	78
11. Working With Hosts	79
11.1. Adding Hosts	79
11.2. Scheduled Maintenance and Maintenance Mode for Hosts	79
11.2.1. vCenter and Maintenance Mode	79
11.2.2. XenServer and Maintenance Mode	79
11.3. Disabling and Enabling Zones, Pods, and Clusters	80
11.4. Removing Hosts	80
11.4.1. Removing XenServer and KVM Hosts	81
11.4.2. Removing vSphere Hosts	81
11.5. Re-Installing Hosts	81
11.6. Mantendo hipervisores em hosts	81
11.7. Changing Host Password	81
11.8. Host Allocation	82
11.8.1. Over-Provisioning and Service Offering Limits	82
11.9. VLAN Provisioning	83
12. Trabalhando com templates	85
12.1. Creating Templates: Overview	85
12.2. Requirements for Templates	85
12.3. Best Practices for Templates	85
12.4. The Default Template	85
12.5. Private and Public Templates	86
12.6. Creating a Template from an Existing Virtual Machine	86
12.7. Creating a Template from a Snapshot	87
12.8. Uploading Templates	87
12.9. Exporting Templates	88
12.10. Creating a Windows Template	89
12.10.1. System Preparation for Windows Server 2008 R2	89
12.10.2. System Preparation for Windows Server 2003 R2	93
12.11. Importing Amazon Machine Images	94
12.12. Converting a Hyper-V VM to a Template	97
12.13. Adding Password Management to Your Templates	98
12.13.1. Instalação do sistema operacional Linux	99
12.13.2. Instalação no Windows	99
12.14. Deleting Templates	99
13. Working With Storage	101
13.1. Storage Overview	101
13.2. Storage primária	101
13.2.1. Best Practices for Primary Storage	101
13.2.2. Runtime Behavior of Primary Storage	101
13.2.3. Hypervisor Support for Primary Storage	101
13.2.4. Storage Tags	102
13.2.5. Maintenance Mode for Primary Storage	102

13.3. Storage secundária	103
13.4. Working With Volumes	103
13.4.1. Creating a New Volume	103
13.4.2. Uploading an Existing Volume to a Virtual Machine	104
13.4.3. Attaching a Volume	105
13.4.4. Detaching and Moving Volumes	105
13.4.5. VM Storage Migration	106
13.4.6. Resizing Volumes	107
13.4.7. Volume Deletion and Garbage Collection	108
13.5. Working with Snapshots	108
13.5.1. Snapshot Job Throttling	109
13.5.2. Automatic Snapshot Creation and Retention	109
13.5.3. Incremental Snapshots and Backup	109
13.5.4. Volume Status	110
13.5.5. Snapshot Restore	110
14. Working with Usage	111
14.1. Configuring the Usage Server	111
14.2. Setting Usage Limits	113
14.3. Globally Configured Limits	113
14.4. Default Account Resource Limits	114
14.5. Per-Domain Limits	115
15. Gerenciando redes e tráfego	117
15.1. Tráfego de hóspedes	117
15.2. Rede em um pod	117
15.3. Rede em uma zona	119
15.4. Configuração de rede física de zona básica	119
15.5. Configuração de rede física de zona avançada	119
15.5.1. Configure o tráfego hóspede na zona avançada	120
15.5.2. Configure o tráfego público na zona avançada	120
15.6. Usando múltiplas redes hóspedes	121
15.6.1. Adicionando uma rede hóspede adicional	121
15.6.2. Alterando a oferta de rede em uma rede hóspede	121
15.7. Grupos de segurança	122
15.7.1. About Security Groups	122
15.7.2. Adicionando um grupo de segurança	122
15.7.3. Security Groups in Advanced Zones (KVM Only)	123
15.7.4. Habilitando grupos de segurança	123
15.7.5. Adicionando regras de ingresso e egresso a um grupo de segurança	123
15.8. Firewalls e balanceadores de carga externos	125
15.8.1. Sobre a utilização do balanceador de carga NetScaler	125
15.8.2. Configuring SNMP Community String on a RHEL Server	126
15.8.3. Configuração inicial de firewalls e balanceadores de carga externos	127
15.8.4. Configuração continuada de firewalls e balanceadores de carga externos	128
15.8.5. Configuring AutoScale	128
15.9. Regras de balanceamento de carga	133
15.9.1. Adding a Load Balancer Rule	133
15.9.2. Sticky Session Policies for Load Balancer Rules	134
15.10. Guest IP Ranges	134
15.11. Obtendo um novo endereço IP	134
15.12. Liberando um endereço IP	135
15.13. NAT estática	135
15.13.1. Habilitando ou desabilitando NAT estática	135
15.14. Encaminhamento de IP e firewall	136

15.14.1. Creating Egress Firewall Rules in an Advanced Zone	136
15.14.2. Regras de firewall	137
15.14.3. Encaminhamento de Porta	138
15.15. Balanceamento de carga de IP	139
15.16. DNS e DHCP	139
15.17. VPN	139
15.17.1. Configurando VPN	140
15.17.2. Usando VPN com Windows	140
15.17.3. Using VPN with Mac OS X	141
15.17.4. Configurando uma conexão VPN Site-to-Site	141
15.18. About Inter-VLAN Routing	148
15.19. Configuring a Virtual Private Cloud	150
15.19.1. About Virtual Private Clouds	150
15.19.2. Adding a Virtual Private Cloud	152
15.19.3. Adding Tiers	152
15.19.4. Configuring Access Control List	154
15.19.5. Adicionando um gateway privado a uma VPC	156
15.19.6. Implantando máquinas virtuais na camada	157
15.19.7. Obtendo um novo endereço IP para uma VPC	157
15.19.8. Liberando um endereço IP atribuído a uma VPC	158
15.19.9. Habilitando ou desabilitando NAT estática em uma VPC	159
15.19.10. Adicionando regras de balanceamento de carga em uma VPC	160
15.19.11. Adicionando uma regra de encaminhamento de porta em uma VPC	161
15.19.12. Removing Tiers	162
15.19.13. Editing, Restarting, and Removing a Virtual Private Cloud	163
15.20. Persistent Networks	163
15.20.1. Persistent Network Considerations	163
15.20.2. Creating a Persistent Guest Network	164
16. Working with System Virtual Machines	165
16.1. The System VM Template	165
16.2. Multiple System VM Support for VMware	165
16.3. Console Proxy	165
16.3.1. Using a SSL Certificate for the Console Proxy	166
16.3.2. Changing the Console Proxy SSL Certificate and Domain	166
16.4. Virtual Router	167
16.4.1. Configuring the Virtual Router	168
16.4.2. Upgrading a Virtual Router with System Service Offerings	168
16.4.3. Best Practices for Virtual Routers	168
16.5. Secondary Storage VM	168
17. System Reliability and High Availability	169
17.1. HA for Management Server	169
17.2. Management Server Load Balancing	169
17.3. HA-Enabled Virtual Machines	169
17.4. HA for Hosts	169
17.4.1. Dedicated HA Hosts	170
17.5. Primary Storage Outage and Data Loss	170
17.6. Secondary Storage Outage and Data Loss	170
17.7. Limiting the Rate of API Requests	171
17.7.1. Configuring the API Request Rate	171
17.7.2. Limitations on API Throttling	171
18. Gerenciando a nuvem	173
18.1. Using Tags to Organize Resources in the Cloud	173

18.2. Changing the Database Configuration	174
18.3. Changing the Database Password	174
18.4. Administrator Alerts	175
18.5. Customizing the Network Domain Name	175
18.6. Stopping and Restarting the Management Server	176
19. Global Configuration Parameters	177
19.1. Setting Global Configuration Parameters	177
19.2. About Global Configuration Parameters	177
20. CloudStack API	181
20.1. Provisioning and Authentication API	181
20.2. Allocators	181
20.3. User Data and Meta Data	181
21. Tuning	183
21.1. Performance Monitoring	183
21.2. Increase Management Server Maximum Memory	183
21.3. Set Database Buffer Pool Size	183
21.4. Set and Monitor Total VM Limits per Host	184
21.5. Configure XenServer dom0 Memory	184
22. Troubleshooting	185
22.1. Events	185
22.1.1. Event Logs	185
22.1.2. Event Notification	185
22.1.3. Standard Events	186
22.1.4. Long Running Job Events	186
22.1.5. Event Log Queries	187
22.2. Working with Server Logs	187
22.3. Data Loss on Exported Primary Storage	188
22.4. Recovering a Lost Virtual Router	188
22.5. Maintenance mode not working on vCenter	189
22.6. Unable to deploy VMs from uploaded vSphere template	189
22.7. Unable to power on virtual machine on VMware	189
22.8. Load balancer rules fail after changing network offering	190
A. Time Zones	191
B. Event Types	193
C. Alerts	195
D. Revision History	197

Conceitos

1.1. O que é o CloudStack?

O CloudStack é uma plataforma de software de código aberto que gerencia recursos computacionais para construir nuvens "Infrastructure as a Service" (IaaS). O CloudStack gerencia a rede, storage, e nós computacionais que compõem a infraestrutura de nuvem. Use o CloudStack para implementar, gerenciar e configurar ambientes de computação em nuvem.

Provedores de serviços e empresas são os usuários típicos. Com o CloudStack, você pode:

- Estabelecer um serviço sob demanda elástico de computação em nuvem. Provedores de serviços podem vender instâncias self service de máquinas virtuais, volumes de armazenamento e configurações de rede pela Internet.
- Estabelecer na empresa uma nuvem privada para uso dos funcionários. Ao invés de gerenciar máquinas virtuais do mesmo modo que máquinas físicas, com o CloudStack uma empresa pode oferecer máquinas virtuais self-service para usuários sem envolver os departamentos de TI.



1.2. O que o CloudStack pode fazer?

Suporte a múltiplos hipervisores

O CloudStack trabalha com uma variedade de hipervisores. Uma única implementação de nuvem pode conter múltiplas implementações de hipervisores. O release atual do CloudStack suporta soluções empresariais pre-packaged como o Citrix XenServer e o VMware vSphere, assim como KVM ou Xen executando no Ubuntu ou CentOS.

Gestão de infraestrutura altamente escalável

O CloudStack pode gerenciar dezenas de milhares de servidores instalados em múltiplos centros de computação geograficamente distribuídos. O servidor de gerenciamento centralizado é linearmente escalável, eliminando a necessidade de servidores de gerenciamento de cluster intermediários. Nenhuma falha de componente único pode causar uma interrupção geral na nuvem. A manutenção periódica do servidor de gerenciamento pode ser executada sem afetar o funcionamento de máquinas virtuais que são executadas na nuvem.

Gerenciamento automático de configuração

O CloudStack automaticamente configura os parâmetros de rede e armazenamento de cada máquina virtual hóspede.

O CloudStack gerencia internamente um pool de dispositivos virtuais ("virtual appliances") para suporte à nuvem. Estes dispositivos oferecem serviços como firewall, roteamento, DHCP, acesso VPN, console proxy, acesso a storage e replicação de storage. O uso extensivo de dispositivos virtuais simplifica a instalação, configuração e gerenciamento contínuo de uma implementação de nuvem.

Interface gráfica do usuário

O CloudStack oferece uma interface web para o administrador, usado para provisionamento e gestão da nuvem, assim como uma interface web do usuário final, usada para executar máquinas virtuais e gerenciar modelos (templates) de máquinas virtuais. A interface de usuário pode ser customizada para refletir os padrões de visuais de apresentação do provedor de serviços ou empresa.

API e extensibilidade

CloudStack provides an API that gives programmatic access to all the management features available in the UI. The API is maintained and documented. This API enables the creation of command line tools and new user interfaces to suit particular needs. See the Developer's Guide and API Reference, both available at [Apache CloudStack Guides](http://cloudstack.apache.org/docs/en-US/index.html)¹ and [Apache CloudStack API Reference](http://cloudstack.apache.org/docs/api/index.html)² respectively.

A arquitetura de alocação de plataformas conectáveis do CloudStack permite a criação de novos tipos de alocadores para a seleção de storage e hosts. Veja o "Allocator Implementation Guide" (http://docs.cloudstack.org/CloudStack_Documentation/Allocator_Implementation_Guide).

Alta disponibilidade

A plataforma do CloudStack tem um número de recursos para aumentar a disponibilidade do sistema. O próprio servidor de gerenciamento pode ser implementado em um ambiente com múltiplos nós onde é feito balanceamento de carga entre os servidores. MySQL pode ser configurado para usar replicação, provendo uma método manual de recuperação em caso de perda do database. Para os hosts, a plataforma CloudStack suporta NIC bonding e o uso de redes isoladas de storage, assim como iSCSI Multipath.


1.3. Visão geral da arquitetura de implementação

Uma instalação do CloudStack consiste em dois componentes: o servidor de gerenciamento e a infraestrutura de nuvem que ele gerencia. Quando você monta e gerencia uma nuvem CloudStack, você provê recursos como hosts, equipamentos de storage, e endereços IP no servidor de gerenciamento, e o servidor de gerenciamento gerencia estes recursos.

¹ <http://cloudstack.apache.org/docs/en-US/index.html>

² <http://cloudstack.apache.org/docs/api/index.html>

A instalação mínima de produção consiste em uma máquina executando o servidor de gerenciamento do CloudStack e outra atuando como a infraestrutura de rede (neste caso, uma infraestrutura muito simples consistindo em um host executando software hipervisor). Na menor implementação possível, uma única máquina pode atuar tanto como servidor de gerenciamento quanto o host hipervisor (usando o hipervisor KVM).



Simplified view of a basic deployment

A more full-featured installation consists of a highly-available multi-node Management Server installation and up to tens of thousands of hosts using any of several advanced networking setups. For information about deployment options, see the "Choosing a Deployment Architecture" section of the \$PRODUCT; Installation Guide.

1.3.1. Visão geral da arquitetura de implementação

O servidor de gerenciamento é o software do CloudStack que gerencia os recursos da nuvem. Pela interação com o servidor de gerenciamento através de sua interface de usuário ou API, você pode configurar e gerenciar sua infraestrutura de nuvem.

O servidor de gerenciamento é executado em um servidor dedicado ou máquina virtual. Ele controla a alocação de máquinas virtuais em hosts e atribui storage e endereços IP às instâncias de máquinas virtuais. O servidor de gerenciamento é executado em um container Tomcat e requer um database MySQL para persistência.

A máquina deve atender os requerimentos descritos em "Requerimentos do sistema".


O servidor de gerenciamento:

- Provê a interface web de usuário para o administrador e uma interface de referência de usuário para usuários finais.
- Provê as APIs para o CloudStack.
- Gerencia o assinalamento de máquinas virtuais hóspedes a hosts determinados.
- Gerencia o assinalamento de endereços IP públicos e privados a determinadas contas.
- Gerencia a alocação de storage - como discos virtuais - para hóspedes.
- Gerencia snapshots, templates, e imagens ISO, possivelmente replicando estes elementos através de data centers.
- Provê um ponto único de configuração para a nuvem.

1.3.2. Visão geral da infraestrutura de nuvem

O servidor de gerenciamento gerencia uma ou mais zonas (tipicamente, datacenters) contendo hosts onde máquinas virtuais hóspedes serão executadas. A infraestrutura de nuvem é organizada como se segue:

- Zona: tipicamente, uma zona é equivalente a um único datacenter. Uma zona consiste em um ou mais pods e storage secundária.
- Pod: um pod é usualmente um rack de hardware que inclui uma switch layer-2 e um ou mais clusters.
- Cluster: um cluster consiste em um ou mais hosts e storage primária.
- Host: um nó computacional em um cluster. É nos hosts onde realmente os serviços de nuvem são executados, na forma de máquinas virtuais hóspedes.
- Storage primária é associada com um cluster, e armazena os volumes de disco para todas as máquinas virtuais sendo executadas em hosts neste cluster.
- Storage secundária é associada com uma zona, e armazena templates, imagens ISO e snapshots de volumes de disco.



Nested organization of a zone

Informações adicionais

Para informações adicionais, veja a documentação sobre conceitos de infraestrutura de nuvem.

1.3.3. Visão geral de serviços de rede

O CloudStack oferece dois tipos de ambiente de rede:

- **Básico.** Para redes no estilo AWS. Provê uma rede única onde isolamento de hóspedes pode ser provido através de recursos da camada 3 como grupos seguros (filtragem de endereço IP de origem).
- **Avançada.** Para topologias de rede mais sofisticadas. Este modelo de rede provê a mais alta flexibilidade na definição de redes hóspedes.

Para mais detalhes, veja Configuração de rede.

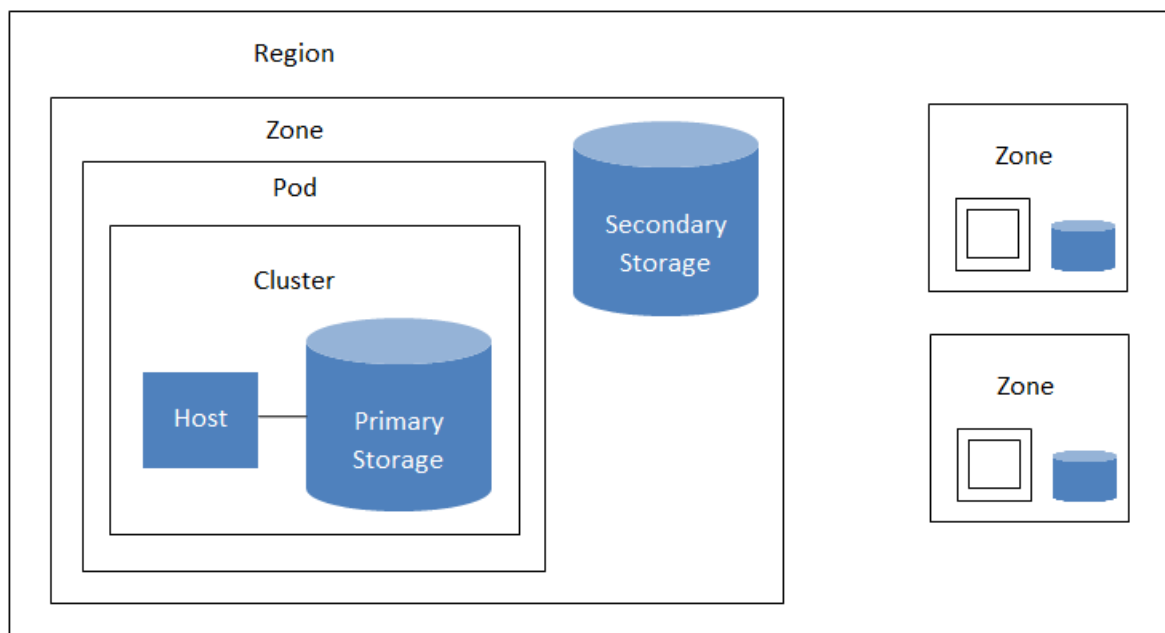
Conceitos de infraestrutura de nuvem

2.1. About Regions

To increase reliability of the cloud, you can optionally group resources into multiple geographic regions. A region is the largest available organizational unit within a CloudStack deployment. A region is made up of several availability zones, where each zone is roughly equivalent to a datacenter. Each region is controlled by its own cluster of Management Servers, running in one of the zones. The zones in a region are typically located in close geographical proximity. Regions are a useful technique for providing fault tolerance and disaster recovery.

By grouping zones into regions, the cloud can achieve higher availability and scalability. User accounts can span regions, so that users can deploy VMs in multiple, widely-dispersed regions. Even if one of the regions becomes unavailable, the services are still available to the end-user through VMs deployed in another region. And by grouping communities of zones under their own nearby Management Servers, the latency of communications within the cloud is reduced compared to managing widely-dispersed zones from a single central Management Server.

Usage records can also be consolidated and tracked at the region level, creating reports or invoices for each geographic region.



A region with multiple zones

Regions are visible to the end user. When a user starts a guest VM, the user must select a region for their guest. Users might also be required to copy their private templates to additional regions to enable creation of guest VMs using their templates in those regions.


2.2. Sobre zonas

A zone is the second largest organizational unit within a CloudStack deployment. A zone typically corresponds to a single datacenter, although it is permissible to have multiple zones in a datacenter. The benefit of organizing infrastructure into zones is to provide physical isolation and redundancy. For

example, each zone can have its own power supply and network uplink, and the zones can be widely separated geographically (though this is not required).

Uma zona consiste em:

- Um ou mais pods. Cada pod contém um ou mais clusters de hosts e um ou mais servidores de storage primária.
- Storage secundária, compartilhada por todos os pods na zona.



Nested organization of a zone

Zonas são visíveis para o usuário final. Quando um usuário ativa uma máquina virtual hóspede, o usuário deve selecionar uma zona para seu hóspede. Usuários podem também ter que copiar seus templates privados para outras zonas para possibilitar a criação de máquinas virtuais hóspedes usando seus templates naquelas zonas.

Zonas podem ser públicas ou privadas. Zonas públicas são visíveis por todos usuários. Isto significa que qualquer usuário pode criar um hóspede na zona. Zonas privadas são reservadas para um domínio específico. Somente usuários no domínio ou seus subdomínios podem criar hóspedes na zona.

Hosts na mesma zona são diretamente acessíveis entre si, sem precisar passar por um firewall. Hosts em zonas distintas podem acessar um ao outro através de túneis VPN configurados estaticamente.

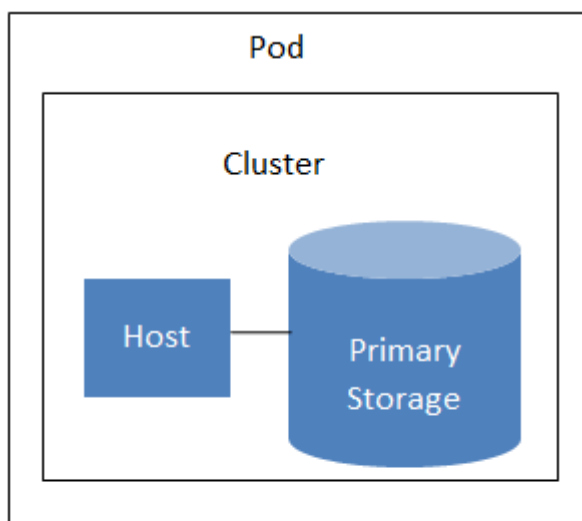
Para cada zona, o administrador deve decidir o seguinte:

- Quantos pods colocar na zona.
- Quantos clusters colocar em cada pod.
- Quantos hosts colocar em em cada cluster.
- Quantos servidores de storage primária colocar em cada cluster e a capacidade total dos servidores de storage.
- Quanto de storage secundária implantar em uma zona.

Quando você adiciona uma zona, você será solicitado a configurar a rede física da zona e adicionar o primeiro pod, primeiro cluster, primeiro host, storage primária inicial, e storage secundária inicial.

2.3. Sobre pods

A pod often represents a single rack. Hosts in the same pod are in the same subnet. A pod is the second-largest organizational unit within a CloudStack deployment. Pods are contained within zones. Each zone can contain one or more pods. A pod consists of one or more clusters of hosts and one or more primary storage servers. Pods are not visible to the end user.




A simple pod

2.4. Sobre clusters

Um cluster provê um modo de agrupar hosts. Mais precisamente, um cluster é um pool de servidores XenServer, um conjunto de servidores KVM, um conjunto de servidores OVM, ou um cluster VMware pré-configurado no vCenter. Todos os hosts em um cluster têm hardwares idênticos, executam o mesmo hipervisor, estão na mesma subnet e acessam a mesma storage primária compartilhada. Instâncias de máquinas virtuais (VMs) ativas em um host podem ser migradas para outro host no mesmo cluster, sem interromper os serviços fornecidos ao usuário.

Um cluster é a terceira maior unidade organizacional em uma implementação do CloudStack. Clusters são contidos em pods, e pods são contidos em zonas. O tamanho do cluster é limitado pelo hipervisor subjacente, embora o CloudStack recomenda menos na maioria dos casos; veja "Melhores práticas".

Um cluster consiste em um ou mais hosts e um ou mais servidores de storage primária.



A simple cluster

O CloudStack permite múltiplos clusters em uma implementação de nuvem.

Mesmo quando exclusivamente armazenamento local é usado, clusters ainda são organizacionalmente requeridos, mesmo que haja somente um host por cluster.

Quando VMware é usado, todo cluster VMware é gerenciado por um servidor vCenter. O administrador deve registrar o servidor vCenter no CloudStack. Pode haver múltiplos servidores vCenter por zona. Cada servidor vCenter pode gerenciar múltiplos clusters VMware.

2.5. Sobre hosts

Um host é um único computador. Hosts fornecem os recursos computacionais que executam as máquinas virtuais hóspedes. Cada host tem software hipervisor instalado para gerenciar as máquinas virtuais hóspedes. Por exemplo, um servidor Linux com KVM habilitado, um servidor Citrix XenServer, e um servidor ESXi são hosts.

Um host é a menor unidade organizacional em uma implementação do CloudStack. Hosts estão contidos em clusters, clusters são contidos em pods, e pods são contidos em zonas.

Hosts em uma implementação do CloudStack:

- Proveem os recursos de CPU, memória, storage e de rede necessários para hospedar as máquinas virtuais
- Interconectam-se utilizando uma rede TCP/IP de alta capacidade e conectam-se à Internet
- Podem residir em múltiplos centros de dados em localidades geograficamente diferentes
- Podem ter diferentes capacidades (diferentes velocidades de CPU, diferentes quantidades de RAM, etc.), embora os hosts em um cluster devam ser homogêneos

Hosts adicionais podem ser adicionados a qualquer momento para prover mais capacidade para máquinas virtuais hóspedes.

O CloudStack automaticamente detecta as quantidades de recursos de CPU e memória fornecidas pelos hosts.

Hosts não são visíveis para o usuário final. Um usuário final não pode determinar a qual host sua máquina hóspede foi assinalada.

Para que um host funcione no CloudStack, você deve fazer o seguinte:

- Instalar software hipervisor no host

- Assinalar um endereço IP ao host
- Certificar-se de que o host está conectado ao servidor de gerenciamento do CloudStack

2.6. Sobre storage primária

Storage primária é associada com um cluster, e armazena os volumes de disco para todas as máquinas virtuais sendo executadas em hosts neste cluster. Você pode adicionar múltiplos servidores de storage primária a um cluster. No mínimo um é requerido. Normalmente é localizado perto dos hosts para um melhor desempenho.

O CloudStack é projetado para trabalhar com todos os servidores compatíveis com os padrões iSCSI e NFS que são suportados pelo hipervisor subjacente, incluindo, por exemplo:

- Dell EqualLogic™ for iSCSI
- Servidores de arquivo Network Appliances for NFS and iSCSI
- Scale Computing for NFS

Se você pretende utilizar somente o disco local na sua instalação, você pode passar para Adição de storage secundária.

2.7. Sobre storage secundária

Storage secundária é associada com um zona, e armazena o seguinte:

- Templates — imagens de sistemas operacionais que podem ser utilizadas no boot de máquinas virtuais e podem incluir configurações adicionais, tais como aplicativos pré-instalados
- Imagens ISO — imagens de disco contendo dados ou mídia bootável de sistema operacional
- Snapshots de volumes de disco — cópias de dados de máquinas virtuais que podem ser usadas para recuperação de dados ou na criação de novos templates

Os itens em storage secundária NFS da zona estão disponíveis para todos os hosts na zona. O CloudStack gerencia a alocação de discos virtuais hóspedes em equipamentos de storage primária.

Para tornar itens na storage secundária disponíveis para todos os hosts na nuvem, você pode adicionar OpenStack Object Storage (Swift, swift.openstack.org¹) em adição à storage secundária NFS da zona. Ao utilizar Swift, você configura storage Swift para todo o CloudStack, em seguida configura storage NFS secundária para cada zona, como usual. A storage NFS em cada zona atua como uma área intermediária através da qual todos os templates e outros dados da storage secundária passam antes de serem encaminhados para a Swift. A storage Swift atua como um recurso da nuvem, tornando disponíveis para qualquer zona na nuvem templates e outros dados. Não há hierarquia na storage Swift, apenas um container Swift por objeto de storage. Qualquer storage secundária na nuvem pode obter um container da Swift quando necessário. Não é necessário copiar templates e snapshots de uma zona para outra, como seria requerido ao utilizar somente storage NFS na zona. Tudo está disponível em todo lugar.

¹ <http://swift.openstack.org>

2.8. Sobre redes físicas

Em parte, a adição de uma zona é configurar a rede física. Uma ou (em uma zona avançada) mais redes físicas podem ser associada com cada zona. A rede corresponde a uma NIC no host hipervisor. Cada rede física pode transportar um ou mais tipos de tráfego de rede. As escolhas de tipo de tráfego para cada rede variam dependendo se você está criando uma zona com rede básica ou rede avançada.

Uma rede física é o hardware de rede e o cabeamento em uma zona. Uma zona pode ter múltiplas redes físicas. Um administrador pode:

- Adicionar/Remover/Atualizar redes físicas em uma zona
- Configurar VLANs na rede física
- Configurar um nome para que a rede seja reconhecida pelos hipervisores
- Configurar os provedores de serviços (firewalls, balanceadores de carga, etc.) disponíveis em uma rede física
- Configurar os endereços IP implementados, ou trunked, para uma rede física
- Especificar o tipo de tráfego que é transportado na rede física, assim como outras propriedades como a velocidade da rede

2.8.1. Tipos de tráfego de rede da zona básica

Quando rede básica é utilizada, somente pode haver um tipo de rede física na zona. Tal rede física transporta os seguintes tipos de tráfego:

- **Hóspede.** Quando usuários finais executam máquinas virtuais, eles geram tráfego hóspede. As máquinas virtuais hóspedes podem se comunicar através de uma rede que pode ser referida como a rede hóspede. Cada pod em uma zona básica é um domínio de broadcast, e portanto cada pod tem um diferente intervalo IP para a rede hóspede. O administrador pode configurar o intervalo IP para cada pod.
- **Management.** When CloudStack's internal resources communicate with each other, they generate management traffic. This includes communication between hosts, system VMs (VMs used by CloudStack to perform various tasks in the cloud), and any other component that communicates directly with the CloudStack Management Server. You must configure the IP range for the system VMs to use.



Nota

Recomendamos fortemente o uso de placas de rede separadas para o tráfego de gerência e tráfego de hóspedes.

- **Público.** Tráfego público é gerado quando máquinas virtuais na nuvem acessam a Internet. Endereços IP publicamente acessíveis devem ser alocados para esta finalidade. Usuários finais podem usar a interface de usuário do CloudStack para obter estes endereços IP para implementar NAT entre sua rede hóspede e a rede pública, como descrito em Obtendo um novo endereço IP.

- **Storage.** While labeled "storage" this is specifically about secondary storage, and doesn't affect traffic for primary storage. This includes traffic such as VM templates and snapshots, which is sent between the secondary storage VM and secondary storage servers. CloudStack uses a separate Network Interface Controller (NIC) named storage NIC for storage network traffic. Use of a storage NIC that always operates on a high bandwidth network allows fast template and snapshot copying. You must configure the IP range to use for the storage network.

Em uma rede básica, a configuração da rede física é bastante simples. Na maioria dos casos, você precisa somente configurar uma rede hóspede para transportar tráfego que é gerado pelas máquinas virtuais hóspedes. Se você usa um balanceador de carga NetScaler e habilita suas características de balanceamento elástico de IP e de carga (EIP e ELB), você pode também configurar uma rede para transportar tráfego público. O CloudStack apresenta na interface de usuário os passos de configuração de rede necessários quando você adiciona uma nova zona.

2.8.2. Endereços IP de hóspedes na zona básica

Quando rede básica é usada, o CloudStack irá assinalar endereços IP no CIDR do pod para hóspedes naquele pod. O administrador deve adicionar um intervalo IP direto no pod para este propósito. Estes IPs estão na mesma VLAN que os hosts.

2.8.3. Tipos de tráfego de rede da zona avançada

Quando rede avançada é usada, pode haver múltiplas redes físicas na zona. Cada rede física pode transportar um ou mais tipos de tráfego, e você precisa informar ao CloudStack qual o tipo de tráfego que você deseja que cada rede transporte. Os tipos de tráfego em uma zona avançada são:

- **Hóspede.** Quando usuários finais executam máquinas virtuais, eles geram tráfego hóspede. As máquinas virtuais hóspedes comunicam-se através de uma rede que pode ser referida como a rede hóspede. Esta rede pode ser isolada ou compartilhada. Em uma rede hóspede isolada, o administrador precisa reservar intervalos de VLAN para prover isolamento para cada rede de conta do CloudStack (potencialmente um grande número de VLANs). Em uma rede hóspede compartilhada, todas as máquinas virtuais hóspedes compartilham uma única rede.
- **Gerência.** Quando recursos internos do CloudStack comunicam-se entre si, eles geram tráfego de gerência. Isto inclui a comunicação entre hosts, máquinas virtuais do sistema (máquinas virtuais usadas pelo CloudStack para executar várias tarefas na nuvem), e qualquer outro componente que se comunica diretamente com o servidor de gerenciamento do CloudStack. Você deve configurar o intervalo de IP para uso das máquinas virtuais do sistema.
- **Público.** Tráfego público é gerado quando máquinas virtuais na nuvem acessam a Internet. Endereços IP publicamente acessíveis devem ser alocados para esta finalidade. Usuários finais podem usar a interface de usuário do CloudStack para obter estes endereços IP para implementar NAT entre sua rede hóspede e a rede pública, como descrito em Obtendo um novo endereço IP no Guia de Administração.
- **Storage.** While labeled "storage" this is specifically about secondary storage, and doesn't affect traffic for primary storage. This includes traffic such as VM templates and snapshots, which is sent between the secondary storage VM and secondary storage servers. CloudStack uses a separate Network Interface Controller (NIC) named storage NIC for storage network traffic. Use of a storage NIC that always operates on a high bandwidth network allows fast template and snapshot copying. You must configure the IP range to use for the storage network.

Cada um destes tipos de tráfego pode estar em uma rede física separada, ou eles podem ser combinados com certas restrições. Quando você usa o assistente Add Zone na interface de usuário para criar uma nova zona, você é orientado a fazer somente escolhas válidas.

2.8.4. Endereços IP de hóspedes na zona avançada

Quando rede avançada é usada, o administrador pode criar redes adicionais para uso dos hóspedes. Estas redes podem abranger a zona e ficar disponíveis para todas as contas, ou elas podem ser restritas para uma única conta, e neste caso somente a conta específica pode criar hóspedes que se ligam a estas redes. As redes são definidas pela identificação da VLAN, intervalo de IP, e gateway. Se quiser, o administrador pode prover milhares destas redes.

2.8.5. Endereços IP públicos na zona avançada

Quando rede avançada é usada, o administrador pode criar redes adicionais para uso dos hóspedes. Estas redes podem abranger a zona e ficar disponíveis para todas as contas, ou elas podem ser restritas para uma única conta, e neste caso somente a conta específica pode criar hóspedes que se ligam a estas redes. As redes são definidas pela identificação da VLAN, intervalo de IP, e gateway. Se quiser, o administrador pode prover milhares destas redes.

2.8.6. Endereços IP reservados pelo sistema

Em cada zona, você precisa configurar um intervalo de endereços IP reservados para a rede de gerência. Esta rede transporta a comunicação entre o servidor de gerenciamento do CloudStack e várias máquinas virtuais do sistema, tais como as máquinas virtuais de storage secundária, as máquinas virtuais de proxy de console, e DHCP.

Os endereços IP reservados devem ser únicos na nuvem. Você não pode, por exemplo, ter um host em uma zona com o mesmo endereço IP privado que um host em outra zona.

Aos hosts em um pod são assinalados endereços IP privados. Estes são tipicamente endereços da RFC1918. As máquinas virtuais de proxy de console e storage secundária também têm endereços IP privados alocados no CIDR do pod onde são criadas.

Certifique-se de que servidores de computação e servidores de gerenciamento utilizem endereços IP fora do intervalo de IP reservado pelo sistema. Por exemplo, suponha que o intervalo de IP reservado pelo sistema se inicia em 192.168.154.2 e termina em 192.168.154.7. O CloudStack pode usar .2 a .7 para máquinas virtuais do sistema. Isto deixa o restante do CIDR do pod, de .8 a .254, para o servidor de gerenciamento e para os hosts hipervisores.

Em todas as zonas:

Forneça IPs privados para o sistema em cada pod e disponibilize-os no CloudStack.

Para o KVM e o XenServer, o número recomendado de IPs privados por pod é um por host. Se você espera que o pod cresça, adicione agora IPs privados suficientes para permitir o crescimento.

Em uma zona que usa rede avançada:

Para zonas com rede avançada, recomendamos o fornecimento de IPs privados suficientes para o número total de usuários mais o requerido pelas máquinas virtuais de sistema. Tipicamente, cerca de 10 IPs adicionais são requeridos para as máquinas virtuais de sistema. Para informações adicionais sobre máquinas virtuais de sistema, veja Trabalhando com máquinas virtuais de sistema no Guia do administrador.

Quando rede avançada está sendo usada, o número de endereços IP privados disponíveis em cada pod varia dependendo de qual hipervisor está sendo executado nos nós daquele pod. O Citrix XenServer e o KVM usam endereços de enlace local — 169.254.0.0/16 ou fe80::/64 —, o que em teoria provê mais de 65.000 endereços IP privados no bloco de endereçamento. Conforme o pod cresce, isto deveria ser mais que suficiente para qualquer número razoável de hosts, assim como endereços IP para roteadores virtuais hóspedes. O VMWare ESXi, em contraste, usa qualquer

esquema de subnet especificada pelo administrador, e o administrador típico provê apenas 255 IPs por pod. Como estes são compartilhados por máquinas físicas, o roteador virtual hóspede, e outras entidades, é possível esgotar os IPs privados ao incrementar a configuração de um pod cujos nós executam ESXi.

Para garantir uma margem adequada para redimensionar o espaço de IPs privados em um pod ESXi que usa rede avançada, use uma ou ambas técnicas a seguir:

- Especifique um bloco CIDR maior para a subnet. Uma máscara de subnet com um sufixo /20 proverá mais de 4.000 endereços IP.
- Crie múltiplos pods, cada um com sua própria subnet. Por exemplo, se você criar 10 pods e cada pod tem 255 IPs, isto irá prover 2.550 endereços IP.

Accounts

3.1. Accounts, Users, and Domains

Contas

An account typically represents a customer of the service provider or a department in a large organization. Multiple users can exist in an account.

Domínios

Accounts are grouped by domains. Domains usually contain multiple accounts that have some logical relationship to each other and a set of delegated administrators with some authority over the domain and its subdomains. For example, a service provider with several resellers could create a domain for each reseller.

For each account created, the Cloud installation creates three different types of user accounts: root administrator, domain administrator, and user.

Usuários

Users are like aliases in the account. Users in the same account are not isolated from each other, but they are isolated from users in other accounts. Most installations need not surface the notion of users; they just have one user per account. The same user cannot belong to multiple accounts.

Username is unique in a domain across accounts in that domain. The same username can exist in other domains, including sub-domains. Domain name can repeat only if the full pathname from root is unique. For example, you can create root/d1, as well as root/foo/d1, and root/sales/d1.

Administrators are accounts with special privileges in the system. There may be multiple administrators in the system. Administrators can create or delete other administrators, and change the password for any user in the system.

Domain Administrators

Domain administrators can perform administrative operations for users who belong to that domain. Domain administrators do not have visibility into physical servers or other domains.

Root Administrator

Root administrators have complete access to the system, including managing templates, service offerings, customer care administrators, and domains

The resources belong to the account, not individual users in that account. For example, billing, resource limits, and so on are maintained by the account, not the users. A user can operate on any resource in the account provided the user has privileges for that operation. The privileges are determined by the role.

3.2. Using an LDAP Server for User Authentication

You can use an external LDAP server such as Microsoft Active Directory or ApacheDS to authenticate CloudStack end-users. Just map CloudStack accounts to the corresponding LDAP accounts using a query filter. The query filter is written using the query syntax of the particular LDAP server, and can

include special wildcard characters provided by CloudStack for matching common values such as the user's email address and name. CloudStack will search the external LDAP directory tree starting at a specified base directory and return the distinguished name (DN) and password of the matching user. This information along with the given password is used to authenticate the user..

To set up LDAP authentication in CloudStack, call the CloudStack API command `ldapConfig` and provide the following:

- Hostname or IP address and listening port of the LDAP server
- Base directory and query filter
- Search user DN credentials, which give CloudStack permission to search on the LDAP server
- SSL keystore and password, if SSL is used

3.2.1. Example LDAP Configuration Commands

To understand the examples in this section, you need to know the basic concepts behind calling the CloudStack API, which are explained in the Developer's Guide.

The following shows an example invocation of `ldapConfig` with an ApacheDS LDAP server

```
http://127.0.0.1:8080/client/api?command=ldapConfig&hostname=127.0.0.1&searchbase=ou%3Dtesting%2Co%3Dproject&queryfilter=%28%26%28uid%3D%25u%29%29&binddn=cn%3DJohn+Singh%2Co%3Dtesting%2Co%3Dproject&bindpass=secret&port=10389&ssl=true&truststore=C%3A%2Fcompany%2Finfo%2Ftrusted.ks&truststorepass=secret&response=json&apiKey=YourAPIKey&signature=YourSignatureHash
```

The command must be URL-encoded. Here is the same example without the URL encoding:

```
http://127.0.0.1:8080/client/api?command=ldapConfig
&hostname=127.0.0.1
&searchbase=ou=testing,o=project
&queryfilter=(&(%uid=%u))
&binddn=cn=John+Singh,ou=testing,o=project
&bindpass=secret
&port=10389
&ssl=true
&truststore=C:/company/info/trusted.ks
&truststorepass=secret
&response=json
&apiKey=YourAPIKey&signature=YourSignatureHash
```

The following shows a similar command for Active Directory. Here, the search base is the testing group within a company, and the users are matched up based on email address.

```
http://10.147.29.101:8080/client/api?command=ldapConfig&hostname=10.147.28.250&searchbase=OU%3Dtesting%2CDC%3Dcompany&queryfilter=%28%26%28mail%3D%25e%29%29%20&binddn=CN%3DAdministrator%2COU%3Dtesting%2CDC%3Dcompany&bindpass=1111_aaaa&port=389&response=json&apiKey=YourAPIKey&signature=YourSignatureHash
```

The next few sections explain some of the concepts you will need to know when filling out the `ldapConfig` parameters.

3.2.2. Search Base

An LDAP query is relative to a given node of the LDAP directory tree, called the search base. The search base is the distinguished name (DN) of a level of the directory tree below which all users can

be found. The users can be in the immediate base directory or in some subdirectory. The search base may be equivalent to the organization, group, or domain name. The syntax for writing a DN varies depending on which LDAP server you are using. A full discussion of distinguished names is outside the scope of our documentation. The following table shows some examples of search bases to find users in the testing department.

LDAP Server	Example Search Base DN
ApacheDS	ou=testing,o=project
Active Directory	OU=testing, DC=company

3.2.3. Query Filter

The query filter is used to find a mapped user in the external LDAP server. The query filter should uniquely map the CloudStack user to LDAP user for a meaningful authentication. For more information about query filter syntax, consult the documentation for your LDAP server.

The CloudStack query filter wildcards are:

Query Filter Wildcard	Description
%u	User name
%e	Email address
%n	First and last name

The following examples assume you are using Active Directory, and refer to user attributes from the Active Directory schema.

If the CloudStack user name is the same as the LDAP user ID:

```
(uid=%u)
```

If the CloudStack user name is the LDAP display name:

```
(displayName=%u)
```

To find a user by email address:

```
(mail=%e)
```

3.2.4. Search User Bind DN

The bind DN is the user on the external LDAP server permitted to search the LDAP directory within the defined search base. When the DN is returned, the DN and passed password are used to authenticate the CloudStack user with an LDAP bind. A full discussion of bind DN is outside the scope of our documentation. The following table shows some examples of bind DN.

LDAP Server	Example Bind DN
ApacheDS	cn=Administrator,dc=testing,ou=project,ou=org
Active Directory	CN=Administrator, OU=testing, DC=company, DC=com

3.2.5. SSL Keystore Path and Password

If the LDAP server requires SSL, you need to enable it in the `ldapConfig` command by setting the parameters `ssl`, `truststore`, and `truststorepass`. Before enabling SSL for `ldapConfig`, you need to get the certificate which the LDAP server is using and add it to a trusted keystore. You will need to know the path to the keystore and the password.

User Services Overview

In addition to the physical and logical infrastructure of your cloud, and the CloudStack software and servers, you also need a layer of user services so that people can actually make use of the cloud. This means not just a user UI, but a set of options and resources that users can choose from, such as templates for creating virtual machines, disk storage, and more. If you are running a commercial service, you will be keeping track of what services and resources users are consuming and charging them for that usage. Even if you do not charge anything for people to use your cloud – say, if the users are strictly internal to your organization, or just friends who are sharing your cloud – you can still keep track of what services they use and how much of them.

4.1. Service Offerings, Disk Offerings, Network Offerings, and Templates

A user creating a new instance can make a variety of choices about its characteristics and capabilities. CloudStack provides several ways to present users with choices when creating a new instance:

- Service Offerings, defined by the CloudStack administrator, provide a choice of CPU speed, number of CPUs, RAM size, tags on the root disk, and other choices. See [Creating a New Compute Offering](#).
- Disk Offerings, defined by the CloudStack administrator, provide a choice of disk size for primary data storage. See [Creating a New Disk Offering](#).
- Network Offerings, defined by the CloudStack administrator, describe the feature set that is available to end users from the virtual router or external networking devices on a given guest network. See [Network Offerings](#).
- Templates, defined by the CloudStack administrator or by any CloudStack user, are the base OS images that the user can choose from when creating a new instance. For example, CloudStack includes CentOS as a template. See [Working with Templates](#).

In addition to these choices that are provided for users, there is another type of service offering which is available only to the CloudStack root administrator, and is used for configuring virtual infrastructure resources. For more information, see [Upgrading a Virtual Router with System Service Offerings](#).

Interface do usuário

5.1. Login na interface de usuário

O CloudStack provê uma interface de usuário web que pode ser usada tanto por administradores e usuários finais. A versão apropriada da interface é apresentada dependendo das credenciais utilizadas para login. A interface de usuário está disponível em navegadores populares, incluindo o IE7, IE8, IE9, Firefox 3.5+, Firefox 4, Safari 4 e Safari 5. A URL é: (substitua pelo endereço IP de seu servidor de gerenciamento)

```
http://<management-server-ip-address>:8080/client
```

Em uma nova instalação do servidor de gerenciamento, uma tela de apresentação é exibida. Em visitas posteriores, você verá uma tela de login onde se especifica o seguinte para seguir para o painel de instrumentos:

Nome de usuário

A identificação de usuário de sua conta. O nome default de usuário é admin.

Senha

A senha associada com a identificação do usuário. A senha para identificação default é password.

Domínio

Se você é um usuário root, deixe este campo em branco.

Se você é um usuário nos subdomínios, forneça o caminho completo para o domínio, excluindo o domínio root.

Por exemplo, suponha que múltiplos níveis são criados sob o domínio root, tal como Comp1/hr. Os usuários no domínio Comp1 devem informar Comp1 no campo Domain, enquanto os usuários no domínio Comp1/sales devem informar Comp1/sales.

Para mais orientação sobre as escolhas apresentadas quando você faz login na interface de usuário, veja [Seção 5.1.3, "Fazendo login como o administrador root"](#).

5.1.1. End User's UI Overview

The CloudStack UI helps users of cloud infrastructure to view and use their cloud resources, including virtual machines, templates and ISOs, data volumes and snapshots, guest networks, and IP addresses. If the user is a member or administrator of one or more CloudStack projects, the UI can provide a project-oriented view.

5.1.2. Root Administrator's UI Overview

The CloudStack UI helps the CloudStack administrator provision, view, and manage the cloud infrastructure, domains, user accounts, projects, and configuration settings. The first time you start the UI after a fresh Management Server installation, you can choose to follow a guided tour to provision your cloud infrastructure. On subsequent logins, the dashboard of the logged-in user appears. The various links in this screen and the navigation bar on the left provide access to a variety of

administrative functions. The root administrator can also use the UI to perform all the same tasks that are present in the end-user's UI.

5.1.3. Fazendo login como o administrador root

Depois do software do servidor de gerenciamento estar instalado e executando, você pode executar a interface de usuário do CloudStack. Esta interface de usuário está disponível para ajudá-lo a montar, visualizar e gerenciar sua infraestrutura de nuvem.

1. Abra o seu navegador web favorito e acesse esta URL. Substitua o endereço IP de seu servidor de gerenciamento:

```
http://<management-server-ip-address>:8080/client
```

Após fazer login em um nova instalação do servidor de gerenciamento, uma tela de apresentação é exibida. Em visitas posteriores, você será levado diretamente ao painel de controle.

2. Se você vê a tela apresentada somente na primeira vez, escolha uma das opções seguintes.
 - **Continue with basic setup.** Escolha esta opção se você está apenas experimentando o CloudStack, e você deseja uma explicação guiada sobre a configuração mais simples possível de modo a iniciar o uso imediatamente. Iremos ajudá-lo a configurar uma nuvem com os seguintes recursos: uma máquina única executando o CloudStack e que provê storage através do NFS; uma máquina única executando máquinas virtuais sob um hipervisor XenServer ou KVM; uma rede pública compartilhada.

Os prompts nesta visita guiada devem lhe fornecer toda a informação que você precisa, mas se você deseja um pouco mais de detalhe, você pode seguir o Guia de instalação de avaliação.

- **I have used CloudStack before.** Escolha esta opção se você já passou pela fase de design e planejou uma implementação mais sofisticada, ou você está pronto para incrementar uma implementação de avaliação que você preparou anteriormente através das telas de configuração básica. Na interface de usuário do administrador, você pode iniciar a utilização dos recursos mais poderosos de CloudPlatform, tais como configuração avançada de rede VLAN, alta disponibilidade, elementos adicionais de rede tais como balanceadores de carga e firewalls, e suporte para múltiplos hipervisores, incluindo Citrix XenServer, KVM, e VMware vSphere.

O painel de instrumentos do administrador root é exibido.

3. Você deve escolher uma nova senha de root. Se você escolheu configuração básica, você será requisitado a criar uma nova senha imediatamente. Se você selecionou usuário experiente, siga os passos em [Seção 5.1.4, "Changing the Root Password"](#).



Atenção


Você está fazendo login como o administrador root. Esta conta gerencia a implementação do CloudStack, incluindo a infraestrutura física. O administrador root pode modificar opções de configuração para alterar funcionalidade básica, criar ou apagar contas de usuários, e tomar várias ações que somente devem ser executadas por uma pessoa autorizada. Por favor, altere a senha default para uma nova e única senha.

5.1.4. Changing the Root Password

During installation and ongoing cloud administration, you will need to log in to the UI as the root administrator. The root administrator account manages the CloudStack deployment, including physical infrastructure. The root administrator can modify configuration settings to change basic functionality, create or delete user accounts, and take many actions that should be performed only by an authorized person. When first installing CloudStack, be sure to change the default password to a new, unique value.

1. Open your favorite Web browser and go to this URL. Substitute the IP address of your own Management Server:

```
http://<management-server-ip-address>:8080/client
```

2. Log in to the UI using the current root user ID and password. The default is admin, password.
3. Click Accounts.
4. Click the admin account name.
5. Click View Users.
6. Click the admin user name.
7. Click the Change Password button. 
8. Type the new password, and click OK.

5.2. Usando as chaves SSH para autenticação.

Além da autenticação de usuário e senha, o CloudStack suporta o uso de chaves SSH para efetuar login em infraestrutura de nuvem para segurança adicional. Você pode usar a API `createSSHKeyPair` para gerar as chaves SSH.

Como cada usuário da nuvem tem sua própria chave SSH, um usuário da nuvem não pode efetuar login em instâncias de outro usuário da nuvem, a menos que eles compartilham seus arquivos de chave SSH. Usando um único par chave SSH, você pode gerenciar várias instâncias.

5.2.1. Criando um template de instância que suporta chaves SSH

Crie um template de instância que suporta chaves SSH.

1. Crie uma nova instância usando o template fornecido pelo cloudstack.

Para informações adicionais na criação de nova instância, veja

2. Faça download do script cloudstack de [The SSH Key Gen Script](http://sourceforge.net/projects/cloudstack/files/SSH%20Key%20Gen%20Script/)¹ para a instância que você criou.

```
wget http://downloads.sourceforge.net/project/cloudstack/SSH%20Key%20Gen%20Script/cloud-set-guest-sshkey.in?r=http%3A%2F%2Fsourceforge.net%2Fprojects%2Fcloudstack%2Ffiles%2FSSH%20Key%20Gen%20Script%2F&ts=1331225219&use_mirror=iweb
```

3. Copie o arquivo para `/etc/init.d`.

¹ <http://sourceforge.net/projects/cloudstack/files/SSH%20Key%20Gen%20Script/>

```
cp cloud-set-guest-sshkey.in /etc/init.d/
```

4. Dê as permissões necessárias ao script:

```
chmod +x /etc/init.d/cloud-set-guest-sshkey.in
```

5. Execute o script ao iniciar o sistema operacional:

```
chkconfig --add cloud-set-guest-sshkey.in
```

6. Pare a instância.

5.2.2. Criando o par de chaves SSH

Você deve fazer uma chamada para o método api createSSHKeyPair. Você pode usar a biblioteca API Python do CloudStack ou os comandos curl para fazer a chamada para a API CloudStack.

Por exemplo, faça uma chamada a partir do servidor CloudStack para criar um par de chaves SSH chamado de "keypair-doc" para a conta admin no domínio root:



Nota

Certifique-se de ajustar esses valores para atender o que você precisa. Se você está fazendo a chamada da API de um servidor diferente, o URL/PORT será diferente, e você vai precisar usar as chaves de API.

1. Execute o seguinte comando curl:

```
curl --globoff "http://localhost:8096/?command=createSSHKeyPair&name=keypair-doc&account=admin&domainid=5163440e-c44b-42b5-9109-ad75cae8e8a2"
```

A saída é algo semelhante ao que é dado a seguir:

```
<?xml version="1.0" encoding="ISO-8859-1"?><createsshkeypairresponse
  cloud-stack-version="3.0.0.20120228045507"><keypair><name>keypair-
doc</name><fingerprint>f6:77:39:d5:5e:77:02:22:6a:d8:7f:ce:ab:cd:b3:56</
fingerprint><privatekey>-----BEGIN RSA PRIVATE KEY-----
MIICXQIBAAKBgQCSydmnQ67jP6lNoXdX3noZjQdrMAWNQZ7y5SrEu4wDxplvhYci
dXYBeZVwakDVsU2MLG1/K+wefwefwefwefwefJyKJaogMKn7BperPD6n1wIDAQAB
AoGAdXaJ7uyZKeRDoy6wA0UmF0kSPbMZCR+UTIHNkS/E0/4U+6lhMokmFShtu
mFDZ1kGGDYhMsdytjDBztljawfawfeawefawfawfawQQDCjEsoRdgkduTy
QpbSGDIa11Jsc+XNDx2fgRinDsxxI/zJYXTKRhSl/LIPHBw/brw8vzxh0lSOrwm7
VvemkkgpAKEAwSeEw394LYZiEVv395ar9MLRVTVLwpo54jC4tsOxQCB1loocK
lYaocpk0yBqq0USBawfIiDCuLXSdvBo1Xz5ICTM19vgvEp/+kMuECQBzm
nVo8b2Gvyagqt/KEQo8wzH2THghZ1qQ1QRhIeJG2aissEacF6bGB2oZ7Igim5L14
4KR70eEToyCLC2k+02UCQQCrniSnWkTDVoVqEK/zbB32JhW3Wu1lv5p5zUEcd
KfEEuzccUIxtJYTahJ1pv1FkQ8anpuxjSEDP8x/18bq3
-----END RSA PRIVATE KEY-----
</privatekey></keypair></createsshkeypairresponse>
```

2. Copie os dados da chave em um arquivo. O arquivo fica assim:

```

-----BEGIN RSA PRIVATE KEY-----
MIICXQIBAAKBgQCsydmnQ67jP6lNoXdX3noZjQdrMAWNQZ7y5SrEu4wDxp1vhYci
dXYBeZVwakDVsu2MLG1/K+wefwefwefwefwefJyKJaogMKn7BperPD6n1wIDAQAB
AoGAdXaJ7uyZKeRDoy6wA0UmF0kSPbMZCR+UTIHnks/E0/4U+6lhMokmFShtu
mfDZ1kGGDYhMsdytjDBzt1jawfawfeawefawfawfawQQDCjEsoRdgkduTy
QpbSGDIa11Jsc+XNDx2fgRinDsxxI/zJYXTKRhSl/LIPHBw/brW8vzxh0lSOrwm7
VvemkkgpAkEAWSeEw394LYZiEVv395ar9MLRVTVLwpo54jC4ts0xQCB1loocK
lYaocpk0yBqq0USBawfIiDCuLXSDvBo1Xz5ICTM19vgvEp/+kMuECQBzm
nVo8b2Gvyagqt/KEQo8wzh2THghZ1qQ1QRhIeJG2aissEacF6bGB2oZ7Igm5L14
4KR70eEToyCLC2k+02UCQQRniSnWktDVoVqeK/zB32Jhw3Wu1lv5p5zUEcd
KfEEuzccUIxtJYtahJ1pv1FkQ8anpuxjSEdp8x/18bq3
-----END RSA PRIVATE KEY-----

```

3. Salve o arquivo.

5.2.3. Criando uma instância

Depois de salvar o arquivo de par de chaves SSH, é necessário criar uma instância usando o template que você criou no [Seção 5.2.1, “Criando um template de instância que suporta chaves SSH”](#). Certifique-se de usar o nome da mesma chave SSH que você criou no [Seção 5.2.2, “Criando o par de chaves SSH”](#).



Nota

Você não pode criar a instância usando o GUI neste momento e associar a instância com o recém-criado par de chaves SSH.

Uma amostra do comando curl para criar uma nova instância é:

```

curl --globoff http://localhost:<port number>/?command=deployVirtualMachine
\&zoneId=1\&serviceOfferingId=18727021-7556-4110-9322-d625b52e0813\&templateId=e899c18a-
ce13-4bbf-98a9-625c5026e0b5\&securitygroupids=ff03f02f-9e3b-48f8-834d-91b822da40c5\&account=admin
\&domainid=1\&keypair=keypair-doc

```

Substitua o template, oferta de serviços e IDs de grupo de segurança (se você estiver usando o recurso de grupo de segurança) que estão em seu ambiente de nuvem.

5.2.4. Fazendo login usando o par de chaves SSH

Para testar se as chaves SSH foram geradas com sucesso, verifique se você pode fazer login para configuração da nuvem.

Por exemplo, de um sistema operacional Linux, execute:

```
ssh -i ~/.ssh/keypair-doc <ip address>
```

O parâmetro `-i` informa ao cliente ssh para usar uma chave ssh encontrada em `~/.ssh/keypair-doc`.

5.2.5. Resetting SSH Keys

With the API command `resetSSHKeyForVirtualMachine`, a user can set or reset the SSH keypair assigned to a virtual machine. A lost or compromised SSH keypair can be changed, and the user

can access the VM by using the new keypair. Just create or register a new keypair, then call `resetSSHKeyForVirtualMachine`.

Using Projects to Organize Users and Resources

6.1. Overview of Projects

Projects are used to organize people and resources. CloudStack users within a single domain can group themselves into project teams so they can collaborate and share virtual resources such as VMs, snapshots, templates, data disks, and IP addresses. CloudStack tracks resource usage per project as well as per user, so the usage can be billed to either a user account or a project. For example, a private cloud within a software company might have all members of the QA department assigned to one project, so the company can track the resources used in testing while the project members can more easily isolate their efforts from other users of the same cloud

You can configure CloudStack to allow any user to create a new project, or you can restrict that ability to just CloudStack administrators. Once you have created a project, you become that project's administrator, and you can add others within your domain to the project. CloudStack can be set up either so that you can add people directly to a project, or so that you have to send an invitation which the recipient must accept. Project members can view and manage all virtual resources created by anyone in the project (for example, share VMs). A user can be a member of any number of projects and can switch views in the CloudStack UI to show only project-related information, such as project VMs, fellow project members, project-related alerts, and so on.

The project administrator can pass on the role to another project member. The project administrator can also add more members, remove members from the project, set new resource limits (as long as they are below the global defaults set by the CloudStack administrator), and delete the project. When the administrator removes a member from the project, resources created by that user, such as VM instances, remain with the project. This brings us to the subject of resource ownership and which resources can be used by a project.


Resources created within a project are owned by the project, not by any particular CloudStack account, and they can be used only within the project. A user who belongs to one or more projects can still create resources outside of those projects, and those resources belong to the user's account; they will not be counted against the project's usage or resource limits. You can create project-level networks to isolate traffic within the project and provide network services such as port forwarding, load balancing, VPN, and static NAT. A project can also make use of certain types of resources from outside the project, if those resources are shared. For example, a shared network or public template is available to any project in the domain. A project can get access to a private template if the template's owner will grant permission. A project can use any service offering or disk offering available in its domain; however, you can not create private service and disk offerings at the project level..

6.2. Configuring Projects

Before CloudStack users start using projects, the CloudStack administrator must set up various systems to support them, including membership invitations, limits on project resources, and controls on who can create projects.

6.2.1. Setting Up Invitations

CloudStack can be set up either so that project administrators can add people directly to a project, or so that it is necessary to send an invitation which the recipient must accept. The invitation can be sent by email or through the user's CloudStack account. If you want administrators to use invitations to add members to projects, turn on and set up the invitations feature in CloudStack.

1. Log in as administrator to the CloudStack UI.
2. Na barra de navegação à esquerda, clique em Global Settings.
3. In the search box, type project and click the search button. 
4. In the search results, you can see a few other parameters you need to set to control how invitations behave. The table below shows global configuration parameters related to project invitations. Click the edit button to set each parameter.

Configuration Parameters	Descrição
project.invite.required	Set to true to turn on the invitations feature.
project.email.sender	The email address to show in the From field of invitation emails.
project.invite.timeout	Amount of time to allow for a new member to respond to the invitation.
project.smtp.host	Name of the host that acts as an email server to handle invitations.
project.smtp.password	(Optional) Password required by the SMTP server. You must also set project.smtp.username and set project.smtp.useAuth to true.
project.smtp.port	SMTP server's listening port.
project.smtp.useAuth	Set to true if the SMTP server requires a username and password.
project.smtp.username	(Optional) User name required by the SMTP server for authentication. You must also set project.smtp.password and set project.smtp.useAuth to true..

5. Restart the Management Server:

```
service cloudstack-management restart
```

6.2.2. Setting Resource Limits for Projects

The CloudStack administrator can set global default limits to control the amount of resources that can be owned by each project in the cloud. This serves to prevent uncontrolled usage of resources such as snapshots, IP addresses, and virtual machine instances. Domain administrators can override these resource limits for individual projects with their domains, as long as the new limits are below the global defaults set by the CloudStack root administrator. The root administrator can also set lower resource limits for any project in the cloud

6.2.2.1. Setting Per-Project Resource Limits


The CloudStack root administrator or the domain administrator of the domain where the project resides can set new resource limits for an individual project. The project owner can set resource limits only if the owner is also a domain or root administrator.

The new limits must be below the global default limits set by the CloudStack administrator (as described in [Seção 6.2.2, "Setting Resource Limits for Projects"](#)). If the project already owns more of

a given type of resource than the new maximum, the resources are not affected; however, the project can not add any new resources of that type until the total drops below the new limit.

1. Log in as administrator to the CloudStack UI.
2. In the left navigation, click Projects.
3. In Select View, choose Projects.
4. Click the name of the project you want to work with.
5. Click the Resources tab. This tab lists the current maximum amount that the project is allowed to own for each type of resource.
6. Type new values for one or more resources.
7. Click Apply.

6.2.2.2. Setting the Global Project Resource Limits

1. Log in as administrator to the CloudStack UI.
2. Na barra de navegação à esquerda, clique em Global Settings.
3. In the search box, type max.projects and click the search button.
4. In the search results, you will see the parameters you can use to set per-project maximum resource amounts that apply to all projects in the cloud. No project can have more resources, but an individual project can have lower limits. Click the edit button to set each parameter. 


max.project.public.ips	Maximum number of public IP addresses that can be owned by any project in the cloud. See About Public IP Addresses.
max.project.snapshots	Maximum number of snapshots that can be owned by any project in the cloud. See Working with Snapshots.
max.project.templates	Maximum number of templates that can be owned by any project in the cloud. See Working with Templates.
max.project.uservms	Maximum number of guest virtual machines that can be owned by any project in the cloud. See Working With Virtual Machines.
max.project.volumes	Maximum number of data volumes that can be owned by any project in the cloud. See Working with Volumes.

5. Reinicie o servidor de gerenciamento.

```
# service cloudstack-management restart
```

6.2.3. Setting Project Creator Permissions

You can configure CloudStack to allow any user to create a new project, or you can restrict that ability to just CloudStack administrators.

1. Log in as administrator to the CloudStack UI.
2. Na barra de navegação à esquerda, clique em Global Settings.
3. In the search box, type `allow.user.create.projects`.
4. Click the edit button to set the parameter. 

<code>allow.user.create.projects</code>	Set to true to allow end users to create projects. Set to false if you want only the CloudStack root administrator and domain administrators to create projects.
---	--

5. Reinicie o servidor de gerenciamento.

```
# service cloudstack-management restart
```

6.3. Creating a New Project

CloudStack administrators and domain administrators can create projects. If the global configuration parameter `allow.user.create.projects` is set to true, end users can also create projects.

1. Log in as administrator to the CloudStack UI.
2. In the left navigation, click Projects.
3. In Select view, click Projects.
4. Click New Project.
5. Give the project a name and description for display to users, then click Create Project.
6. A screen appears where you can immediately add more members to the project. This is optional. Click Next when you are ready to move on.
7. Click Save.

6.4. Adding Members to a Project

New members can be added to a project by the project's administrator, the domain administrator of the domain where the project resides or any parent domain, or the CloudStack root administrator. There are two ways to add members in CloudStack, but only one way is enabled at a time:

- If invitations have been enabled, you can send invitations to new members.
- If invitations are not enabled, you can add members directly through the UI.

6.4.1. Sending Project Membership Invitations

Use these steps to add a new member to a project if the invitations feature is enabled in the cloud as described in [Seção 6.2.1, "Setting Up Invitations"](#). If the invitations feature is not turned on, use the procedure in Adding Project Members From the UI.

1. Log in to the CloudStack UI.
2. In the left navigation, click Projects.

3. In Select View, choose Projects.
4. Click the name of the project you want to work with.
5. Click the Invitations tab.
6. In Add by, select one of the following:
 - a. Account – The invitation will appear in the user's Invitations tab in the Project View. See [Using the Project View](#).
 - b. Email – The invitation will be sent to the user's email address. Each emailed invitation includes a unique code called a token which the recipient will provide back to CloudStack when accepting the invitation. Email invitations will work only if the global parameters related to the SMTP server have been set. See [Seção 6.2.1, "Setting Up Invitations"](#).
7. Type the user name or email address of the new member you want to add, and click Invite. Type the CloudStack user name if you chose Account in the previous step. If you chose Email, type the email address. You can invite only people who have an account in this cloud within the same domain as the project. However, you can send the invitation to any email address.
8. To view and manage the invitations you have sent, return to this tab. When an invitation is accepted, the new member will appear in the project's Accounts tab.

6.4.2. Adding Project Members From the UI

The steps below tell how to add a new member to a project if the invitations feature is not enabled in the cloud. If the invitations feature is enabled cloud, as described in [Seção 6.2.1, "Setting Up Invitations"](#), use the procedure in [Seção 6.4.1, "Sending Project Membership Invitations"](#).

1. Log in to the CloudStack UI.
2. In the left navigation, click Projects.
3. In Select View, choose Projects.
4. Click the name of the project you want to work with.
5. Click the Accounts tab. The current members of the project are listed.
6. Type the account name of the new member you want to add, and click Add Account. You can add only people who have an account in this cloud and within the same domain as the project.

6.5. Accepting a Membership Invitation

If you have received an invitation to join a CloudStack project, and you want to accept the invitation, follow these steps:

1. Log in to the CloudStack UI.
2. In the left navigation, click Projects.
3. In Select View, choose Invitations.
4. If you see the invitation listed onscreen, click the Accept button.

Invitations listed on screen were sent to you using your CloudStack account name.

5. If you received an email invitation, click the Enter Token button, and provide the project ID and unique ID code (token) from the email.


6.6. Suspending or Deleting a Project


When a project is suspended, it retains the resources it owns, but they can no longer be used. No new resources or members can be added to a suspended project.

When a project is deleted, its resources are destroyed, and member accounts are removed from the project. The project's status is shown as Disabled pending final deletion.

A project can be suspended or deleted by the project administrator, the domain administrator of the domain the project belongs to or of its parent domain, or the CloudStack root administrator.

1. Log in to the CloudStack UI.
2. In the left navigation, click Projects.
3. In Select View, choose Projects.
4. Click the name of the project.
5. Click one of the buttons:

To delete, use 

To suspend, use 

6.7. Using the Project View

If you are a member of a project, you can use CloudStack's project view to see project members, resources consumed, and more. The project view shows only information related to one project. It is a useful way to filter out other information so you can concentrate on a project status and resources.

1. Log in to the CloudStack UI.
2. Click Project View.
3. The project dashboard appears, showing the project's VMs, volumes, users, events, network settings, and more. From the dashboard, you can:
 - Click the Accounts tab to view and manage project members. If you are the project administrator, you can add new members, remove members, or change the role of a member from user to admin. Only one member at a time can have the admin role, so if you set another user's role to admin, your role will change to regular user.
 - (If invitations are enabled) Click the Invitations tab to view and manage invitations that have been sent to new project members but not yet accepted. Pending invitations will remain in this list until the new member accepts, the invitation timeout is reached, or you cancel the invitation.

Passos para provisionamento de sua infraestrutura de nuvem

This section tells how to add regions, zones, pods, clusters, hosts, storage, and networks to your cloud. If you are unfamiliar with these entities, please begin by looking through [Capítulo 2, Conceitos de infraestrutura de nuvem](#).


7.1. Visão geral dos passos de provisionamento

Depois do servidor de gerenciamento ser instalado e estar funcionando, você pode adicionar os recursos de computação para gerenciá-lo. Para uma visão geral de como uma infraestrutura de nuvem CloudStack é organizada, veja [Seção 1.3.2, “Visão geral da infraestrutura de nuvem”](#).

Para provisionar a infraestrutura de nuvem, ou para escalá-la a qualquer momento, siga estes procedimentos:

1. Define regions (optional). See [Seção 7.2, “Adding Regions \(optional\)”](#).
2. Add a zone to the region. See [Seção 7.3, “Adicionando uma zona”](#).
3. Add more pods to the zone (optional). See [Seção 7.4, “Adicionando um pod”](#).
4. Add more clusters to the pod (optional). See [Seção 7.5, “Adicionando um cluster”](#).
5. Add more hosts to the cluster (optional). See [Seção 7.6, “Adding a Host”](#).
6. Add primary storage to the cluster. See [Seção 7.7, “Adicionar Storage Primário”](#).
7. Add secondary storage to the zone. See [Seção 7.8, “Adicionar Storage Secundário”](#).
8. Inicialize e teste a nova nuvem. Veja [Seção 7.9, “Initialize and Test”](#).

Quando você terminar estes passos, você terá uma implantação com a seguinte estrutura básica:



Conceptual view of a basic deployment

7.2. Adding Regions (optional)

Grouping your cloud resources into geographic regions is an optional step when provisioning the cloud. For an overview of regions, see [Seção 2.1, “About Regions”](#).

7.2.1. The First Region: The Default Region

If you do not take action to define regions, then all the zones in your cloud will be automatically grouped into a single default region. This region is assigned the region ID of 1.

You can change the name or URL of the default region by using the API command `updateRegion`. For example:

```
http://<IP_of_Management_Server>:8080/client/api?
command=updateRegion&id=1&name=Northern&endpoint=http://
<region_1_IP_address_here>:8080/client&apiKey=miVr6X7u6bN_sdah0BpjNejPgEsT35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU%2FcaiK8RAP001hU
%3D
```

7.2.2. Adding a Region

Use these steps to add a second region in addition to the default region.

1. Each region has its own CloudStack instance. Therefore, the first step of creating a new region is to install the Management Server software, on one or more nodes, in the geographic area where you want to set up the new region. Use the steps in the Installation guide. When you come to the step where you set up the database, use the additional command-line flag `-r <region_id>` to set a region ID for the new region. The default region is automatically assigned a region ID of 1, so your first additional region might be region 2.

```
cloudstack-setup-databases cloud:<dbpassword>@localhost --deploy-as=root:<password> -e
<encryption_type> -m <management_server_key> -k <database_key> -r <region_id>
```

2. By the end of the installation procedure, the Management Server should have been started. Be sure that the Management Server installation was successful and complete.
3. Add region 2 to region 1. Use the API command addRegion. (For information about how to make an API call, see the Developer's Guide.)

```
http://<IP_of_region_1_Management_Server>:8080/client/
api?command=addRegion&id=2&name=Western&endpoint=http://
<region_2_IP_address_here>:8080/client&apiKey=miVr6X7u6bN_sdah0BpjNejPgEst35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU
%2FcaiK8RAP001hU%3D
```

4. Now perform the same command in reverse, adding region 1 to region 2.

```
http://<IP_of_region_2_Management_Server>:8080/client/
api?command=addRegion&id=1&name=Northern&endpoint=http://
<region_1_IP_address_here>:8080/client&apiKey=miVr6X7u6bN_sdah0BpjNejPgEst35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU
%2FcaiK8RAP001hU%3D
```

5. Copy the account, user, and domain tables from the region 1 database to the region 2 database.

In the following commands, it is assumed that you have set the root password on the database, which is a CloudStack recommended best practice. Substitute your own MySQL root password.

- a. First, run this command to copy the contents of the database:

```
# mysqldump -u root -p<mysql_password> -h <region1_db_host> cloud account user domain
> region1.sql
```

- b. Then run this command to put the data onto the region 2 database:

```
# mysql -u root -p<mysql_password> -h <region2_db_host> cloud < region1.sql
```

6. Remove project accounts. Run these commands on the region 2 database:

```
mysql> delete from account where type = 5;
```

7. Set the default zone as null:

```
mysql> update account set default_zone_id = null;
```

8. Restart the Management Servers in region 2.

7.2.3. Adding Third and Subsequent Regions

To add the third region, and subsequent additional regions, the steps are similar to those for adding the second region. However, you must repeat certain steps additional times for each additional region:

Capítulo 7. Passos para provisionamento de sua infraestrutura de nuvem

1. Install CloudStack in each additional region. Set the region ID for each region during the database setup step.

```
cloudstack-setup-databases cloud:<dbpassword>@localhost --deploy-as=root:<password> -e
<encryption_type> -m <management_server_key> -k <database_key> -r <region_id>
```

2. Once the Management Server is running, add your new region to all existing regions by repeatedly calling the API command addRegion. For example, if you were adding region 3:

```
http://<IP_of_region_1_Management_Server>:8080/client/
api?command=addRegion&id=3&name=Eastern&endpoint=http://
<region_3_IP_address_here>:8080/client&apiKey=miVr6X7u6bN_sdah0BpjNejPgEsT35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU
%2FcaiK8RAP001hU%3D
```

```
http://<IP_of_region_2_Management_Server>:8080/client/
api?command=addRegion&id=3&name=Eastern&endpoint=http://
<region_3_IP_address_here>:8080/client&apiKey=miVr6X7u6bN_sdah0BpjNejPgEsT35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU
%2FcaiK8RAP001hU%3D
```

3. Repeat the procedure in reverse to add all existing regions to the new region. For example, for the third region, add the other two existing regions:

```
http://<IP_of_region_3_Management_Server>:8080/client/
api?command=addRegion&id=1&name=Northern&endpoint=http://
<region_1_IP_address_here>:8080/client&apiKey=miVr6X7u6bN_sdah0BpjNejPgEsT35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU
%2FcaiK8RAP001hU%3D
```

```
http://<IP_of_region_3_Management_Server>:8080/client/
api?command=addRegion&id=2&name=Western&endpoint=http://
<region_2_IP_address_here>:8080/client&apiKey=miVr6X7u6bN_sdah0BpjNejPgEsT35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU
%2FcaiK8RAP001hU%3D
```

4. Copy the account, user, and domain tables from any existing region's database to the new region's database.

In the following commands, it is assumed that you have set the root password on the database, which is a CloudStack recommended best practice. Substitute your own MySQL root password.

- a. First, run this command to copy the contents of the database:

```
# mysqldump -u root -p<mysql_password> -h <region1_db_host> cloud account user domain
> region1.sql
```

- b. Then run this command to put the data onto the new region's database. For example, for region 3:

```
# mysql -u root -p<mysql_password> -h <region3_db_host> cloud < region1.sql
```

5. Remove project accounts. Run these commands on the region 2 database:

```
mysql> delete from account where type = 5;
```

- Set the default zone as null:

```
mysql> update account set default_zone_id = null;
```

- Restart the Management Servers in the new region.

7.2.4. Deleting a Region


To delete a region, use the API command `removeRegion`. Repeat the call to remove the region from all other regions. For example, to remove the 3rd region in a three-region cloud:

```
http://<IP_of_region_1_Management_Server>:8080/client/api?
command=removeRegion&id=3&apiKey=miVr6X7u6bN_sdah0BpjNejPgEsT35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU%2FcaiK8RAP001hU
%3D

http://<IP_of_region_2_Management_Server>:8080/client/api?
command=removeRegion&id=3&apiKey=miVr6X7u6bN_sdah0BpjNejPgEsT35eXq-
jB8CG20YI3yaxXcgpyuaIRmFI_EJTVwZ0nUkkJbPmY3y2bciKwFQ&signature=Lxx1DM40AjcXU%2FcaiK8RAP001hU
%3D
```

7.3. Adicionando uma zona

Nestes passos se supõe que você já está logado na interface de usuário do CloudStack. Veja [Seção 5.1, “Login na interface de usuário”](#).

- (Opcional) Se você irá usar Swift para storage secundária na nuvem, você precisa adicioná-la antes de adicionar zonas.
 - Faça login na interface de usuário do CloudStack como administrador.
 - Se esta é a primeira vez utilizando a interface de usuário, você verá uma tela de apresentação. Selecione “Experienced user.” O painel de instrumentos é exibido.
 - Na barra de navegação à esquerda, clique em Global Settings.
 - Na caixa de pesquisa, digite `swift.enable` e clique no botão de pesquisa.
 - Clique no botão de edit e marque `swift.enable` como `true`. 
 - Reinicie o servidor de gerenciamento.

```
# service cloudstack-management restart
```

- Recarregue a interface de usuário do CloudStack no navegador e faça logon novamente.
- Na barra de navegação à esquerda, selecione Infrastructure.
 - Em Zones, clique em View More.
 - (Opcional) Se você está utilizando storage Swift, clique Enable Swift. Forneça o seguinte:
 - URL.** A URL Swift.
 - Account.** A conta Swift.

- **Username.** O nome de usuário associado à conta Swift.
 - **Key.** A chave Swift.
5. Clique em Add Zone. O consultor de criação de zona será exibido.
 6. Selecione um dos seguintes tipos de rede:
 - **Basic.** Para redes no estilo AWS. Provê um rede única a cada instância de máquina virtual é assinalado um endereço IP diretamente da rede. O isolamento de hóspedes pode ser provido através de recursos da camada 3 como grupos seguros (filtragem de endereço IP de origem).
 - **Advanced.** Para topologias de rede mais sofisticadas. Este modelo de rede provê a mais alta flexibilidade na definição de redes hóspedes e oferece customização de rede como firewall, VPN, ou suporte a balanceador de carga.

For more information about the network types, see [Seção 2.8, “Sobre redes físicas”](#).

7. O restante dos passos diferem dependendo se você escolheu Basic ou Advanced. Continue com os passos que se aplicam a você:
 - [Seção 7.3.1, “Configuração de zona básica”](#)
 - [Seção 7.3.2, “Advanced Zone Configuration”](#)

7.3.1. Configuração de zona básica

1. Após selecionar Basic no assistente de Adicionar Zona e clicar em Next, você será solicitado a digitar os seguintes detalhes. Em seguida, clique em Next.
 - **Name.** O nome da zona.
 - **DNS 1 and 2.** Estes são os servidores DNS utilizados pelas máquinas virtuais na zona. Estes servidores serão acessados pela rede publica que será adicionada posteriormente. Os IPs públicos da zona deverão ter acesso a estes servidores.
 - **Internal DNS 1 and Internal DNS 2.** Estes são os servidores DNS utilizados pelas máquinas virtuais na zona (estas são máquinas virtuais usadas pelo CloudStack tais como os roteadores virtuais, proxies console e máquinas virtuais de storage secundária). Estes servidores DNS serão acessados através da interface de gerenciamento de tráfego rede das máquinas virtuais de sistema. O endereço IP privado que você fornecer para os pods devem ter uma rota para o servidor DNS interno identificado aqui.
 - **Hypervisor.** (Introduzido na versão 3.0.1) Escolha o hipervisor para primeiro cluster na zona. Você pode adicionar clusters com diferentes hipervisores mais tarde, depois de você terminar de adicionar a zona.
 - **Network Offering.** Sua escolha aqui determina quais serviços de rede estarão disponíveis na rede para máquinas virtuais hóspedes.

Network Offering	Descrição
DefaultSharedNetworkOfferingWithSGService	Se você quer habilitar grupos de segurança para isolamento de tráfego de hóspedes, escolha esta. (Veja Utilizando grupos de segurança para controlar o tráfego de máquinas virtuais).

Network Offering	Descrição
DefaultSharedNetworkOffering	Se você não precisa de grupos de segurança, escolha esta.
DefaultSharedNetscalerEIPandELBNetworkOffering	Se você tiver instalado o dispositivo Citrix NetScaler como parte de sua rede de zona, e você estará usando o seu IP elástico e características de balanceamento de carga elástica, escolha esta. Com os recursos de EIP e ELB, uma zona básica com grupos de segurança habilitados pode oferecer NAT estático 1:1 e balanceamento de carga.

- **Network Domain.** (Opcional) Se você quiser atribuir um nome de domínio especial à rede de máquina virtual hóspede, especifique o sufixo DNS.
- **Public.** Uma zona pública está disponível para todos usuários. Uma zona que não é pública será atribuída a um domínio específico. Somente a usuários nesse domínio será permitido criar máquinas virtuais hóspedes nesta zona.

2. Escolha os tipos de tráfego que serão transportados pela rede física.

Os tipos de tráfego são: gerência, público, hóspede, e storage. Para mais informações sobre os tipos, role sobre os ícones para exibir suas dicas de ferramentas, ou veja Tipos de tráfego de rede de zona básica. Esta tela começa com alguns tipos de tráfego já atribuídos. Para adicionar mais, arraste e solte os tipos de tráfego na rede. Você também pode alterar o nome da rede, se desejar.

3. 3. (Introduced in version 3.0.1) Assign a network traffic label to each traffic type on the physical network. These labels must match the labels you have already defined on the hypervisor host. To assign each label, click the Edit button under the traffic type icon. A popup dialog appears where you can type the label, then click OK.

Essas etiquetas de tráfego serão definidas apenas para o hipervisor selecionado para o primeiro cluster. Para todos os outros hipervisores, as etiquetas podem ser configuradas depois que a zona é criada.

4. Clique em Next.

5. (Somente NetScaler) Se você escolheu a oferta de rede para NetScaler, você tem uma tela adicional para preencher. Forneça as informações solicitadas para configurar o NetScaler, em seguida, clique em Next.

- **IP address.** O endereço NSIP (NetScaler IP) do dispositivo NetScaler.
- **Username/Password.** A autenticação de credenciais para acessar o dispositivo. O CloudStack usa essas credenciais para acessar o dispositivo.
- **Type.** O tipo de dispositivo NetScaler que está sendo adicionado. Pode ser NetScaler VPX, o NetScaler MPX, ou NetScaler SDX. Para uma comparação dos tipos, veja Sobre como usar um balanceador de carga NetScaler.
- **Public interface.** A Interface de NetScaler que está configurada para ser parte da rede pública.
- **Private interface.** A Interface de NetScaler que está configurada para ser parte da rede privada.

- **Number of retries.** Número de vezes para tentar um comando no dispositivo antes de considerar falhas na operação. O default é 2.
 - **Capacity.** Número de redes hóspedes/contas que irão compartilhar este dispositivo NetScaler.
 - **Dedicated.** Quando marcado como dedicado, este dispositivo será dedicado a uma única conta. Quando dedicado é verificado, o valor no campo Capacity não tem significado - implicitamente, o seu valor é 1.
6. (NetScaler apenas) Configure a faixa de IP para tráfego público. Os IPs nesta faixa serão usados para a capacidade de NAT estática que você ativou selecionando a oferta de rede para NetScaler com EIP e ELB. Insira os detalhes seguintes, e então clique em Add. Se desejar, você pode repetir este passo para adicionar mais intervalos de IP. Quando terminar, clique em Next.
- **Gateway.** O gateway em uso para estes endereços IPs.
 - **Netmask.** A netmask associada com este intervalo de IP.
 - **VLAN.** A VLAN que será usada pelo tráfego público.
 - **Start IP/End IP.** Um intervalo de endereços IP que se supõe serem acessíveis da Internet e que serão alocados para acesso a máquinas virtuais hóspedes.
7. Em uma nova zona, o CloudStack adiciona o primeiro pod para você. Você sempre pode adicionar pods mais tarde. Para uma visão geral do que é pod, veja [Seção 2.3, “Sobre pods”](#).

Para configurar o primeiro pod, informe o que se segue, então clique em Next:

- **Pod Name.** Um nome para o pod.
 - **Reserved system gateway.** O gateway para os hosts no pod.
 - **Reserved system netmask.** O prefixo de rede que define a sub-rede do pod. Use notação CIDR.
 - **Start/End Reserved System IP.** O intervalo IP na rede de gerenciamento que o CloudStack usa para gerenciar várias máquinas virtuais de sistema, tais como as máquinas virtuais de storage secundária, máquinas virtuais de proxy de console, e DHCP. Para mais informações, veja Endereços IP reservados pelo sistema.
8. Configure a rede para o tráfego de hóspedes. Forneça o que se segue, então clique em Next:
- **Guest gateway.** O gateway que os hóspedes devem utilizar.
 - **Guest netmask.** A máscara de rede em uso na sub-rede que os hóspedes utilizarão.
 - **Guest start IP/End IP.** Informe o primeiro e o último endereços IP que definem um intervalo que o CloudStack pode atribuir aos convidados.
 - Recomendamos fortemente o uso de várias placas de rede. Se várias placas de rede são usadas, elas podem estar em uma sub-rede diferente.
 - Se uma placa de rede é utilizada, estes IPs devem estar na mesma CIDR que o CIDR do pod.
9. Em um novo pod, o CloudStack adiciona o primeiro cluster para você. Você sempre pode adicionar clusters mais tarde. Para uma visão geral do que um cluster é, veja Sobre clusters.

Para configurar o primeiro cluster, informe o que segue, então clique em Next:

- **Hypervisor.** (Versão 3.0.0 apenas; na 3.0.1, este campo é somente leitura) Escolha o tipo de software hipervisor que todos os hosts deste cluster executarão. Se você escolher VMware, campos adicionais aparecem para que você possa dar informações sobre um cluster vSphere. Para os servidores vSphere, recomendamos criar o cluster de hosts em vCenter e então adicionar o cluster inteiro no CloudStack. Veja Adicionar cluster: vSphere.
- **Cluster name.** Informe um nome para o cluster. Este pode ser um texto de sua escolha e não é usado pelo CloudStack.

10. Em um novo cluster, o CloudStack adiciona o primeiro host para você. Você sempre pode adicionar mais hosts depois. Para uma visão geral do que um host é, veja Sobre hosts.



Nota

Quando você adicionar um host hipervisor ao CloudStack, o host não deve ter nenhuma máquina virtual já executando.

Antes de configurar o host, você precisa instalar o software hipervisor no host. Você precisará saber qual versão do software hipervisor é suportada pelo CloudStack e qual configuração adicional é requerida para garantir que o host trabalhará com CloudStack. Para encontrar detalhes dessa instalação, veja:

- Instalação e configuração do Citrix XenServer
- Instalação e configuração do VMware vSphere
- Instalação e configuração do KVM

Para configurar o primeiro host, informe o que segue, então clique em Next:

- **Host Name.** O nome DNS ou endereço IP do host.
- **Username.** O nome do usuário é root.
- **Password.** Esta é a senha para o usuário identificado acima (da sua instalação XenServer ou KVM).
- **Host Tags.** (Opcional) Qualquer rótulo que você usa para categorizar hosts para facilidade de manutenção. Por exemplo, você pode definir isso para o rótulo de alta disponibilidade da nuvem (definido no parâmetro de configuração `ha.tag global`) se você quer este host a ser usado apenas para máquinas virtuais com o recurso "alta disponibilidade" habilitado. Para mais informações, veja Máquinas virtuais HA-Enabled, assim como HA para hosts.

11. Em um novo cluster, o CloudStack acrescenta o primeiro servidor de storage primária para você. Você sempre pode adicionar mais servidores mais tarde. Para uma visão geral do que é storage primária, veja Sobre storage primária.

Para configurar o primeiro servidor de storage primária, entre o que segue, então clique em Next:

- **Name.** O nome do dispositivo do storage.

- **Protocol.** Para XenServer, escolha NFS, iSCSI ou PreSetup. Para o KVM, escolha NFS, SharedMountPoint, CLVM, ou RBD. Para vSphere escolha VMFS (iSCSI ou FiberChannel) ou NFS. Os campos restantes na tela variam dependendo do que você escolher aqui.

7.3.2. Advanced Zone Configuration

1. After you select Advanced in the Add Zone wizard and click Next, you will be asked to enter the following details. Then click Next.

- **Nome.** O nome da zona.
- **DNS 1 e 2.** Serão os servidores DNS utilizados pelas máquinas virtuais na Zona. Estes servidores serão acessados pela rede publica que será adicionada a frente. Os ips públicos da zona deverão ter acesso a estes servidores.
- **Internal DNS 1 and Internal DNS 2.** These are DNS servers for use by system VMs in the zone(these are VMs used by CloudStack itself, such as virtual routers, console proxies,and Secondary Storage VMs.) These DNS servers will be accessed via the management traffic network interface of the System VMs. The private IP address you provide for the pods must have a route to the internal DNS server named here.
- **Network Domain.** (Opcional) Se você quiser atribuir um nome de domínio especial à rede de máquina virtual convidada, especificar o sufixo DNS.
- **Guest CIDR.** This is the CIDR that describes the IP addresses in use in the guest virtual networks in this zone. For example, 10.1.1.0/24. As a matter of good practice you should set different CIDRs for different zones. This will make it easier to set up VPNs between networks in different zones.
- **Hypervisor.** (Introduced in version 3.0.1) Escolha o hipervisor para primeiro cluster na zona.Você pode adicionar clusters com diferentes hipervisors mais tarde
- **Public.** Uma zona pública está disponível para todos usuários.Uma zona que não é público será atribuído a um domínio particular. Somente usuários nesse domínio serão permitido criar máquinas virtuais convidadas nesta zona.

2. Escolher os tipos de tráfego serão transmitidos pela rede física.

The traffic types are management, public, guest, and storage traffic. For more information about the types, roll over the icons to display their tool tips, or see [Seção 2.8.3, “Tipos de tráfego de rede da zona avançada”](#). This screen starts out with one network already configured. If you have multiple physical networks, you need to add more. Drag and drop traffic types onto a greyed-out network and it will become active. You can move the traffic icons from one network to another; for example, if the default traffic types shown for Network 1 do not match your actual setup, you can move them down. You can also change the network names if desired.

3. (Introduced in version 3.0.1) Assign a network traffic label to each traffic type on each physical network. These labels must match the labels you have already defined on the hypervisor host. To assign each label, click the Edit button under the traffic type icon within each physical network. A popup dialog appears where you can type the label, then click OK.

Esses label de tráfego serão definidos apenas para o hipervisor selecionado para o primeiro cluster. Para todos os outros hipervisors, os rótulos podem ser configurados depois que a zona é criado.

4. Clicar em Next

5. Configure the IP range for public Internet traffic. Enter the following details, then click Add. If desired, you can repeat this step to add more public Internet IP ranges. When done, click Next.
 - **Gateway.** O gateway em uso para estes endereços IPs.
 - **Netmask.** O netmask associado com esta faixa de IP
 - **VLAN.** A VLAN que será usada pelo tráfego público will be used for public traffic.
 - **Start IP/End IP.** A range of IP addresses that are assumed to be accessible from the Internet and will be allocated for access to guest networks.
6. Em uma nova zona, CloudStack acrescenta primeiro pod para você. Você sempre pode adicionar pods mais tarde. Para uma visão geral do que é pod, veja [Seção 2.3, “Sobre pods”](#).

Para configurar o primeiro pod, enter o seguinte, clique em Next

- **Pod Name.** Um nome para o pod.
 - **Reserved system gateway.** O gateway para hóspedes em que pod
 - **Reserved system netmask.** O prefixo de rede que define a subrede do pod. Use notação CIDR.
 - **Start/End Reserved System IP.** The IP range in the management network that CloudStack uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP. For more information, see [Seção 2.8.6, “Endereços IP reservados pelo sistema”](#).
7. Specify a range of VLAN IDs to carry guest traffic for each physical network (see VLAN Allocation Example), then click Next.
 8. In a new pod, CloudStack adds the first cluster for you. You can always add more clusters later. For an overview of what a cluster is, see [Seção 2.4, “Sobre clusters”](#).

Para configurar o primeiro cluster, Entre o seguinte, então clique Next:

- **Hypervisor.** (Version 3.0.0 only; in 3.0.1, this field is read only) Choose the type of hypervisor software that all hosts in this cluster will run. If you choose VMware, additional fields appear so you can give information about a vSphere cluster. For vSphere servers, we recommend creating the cluster of hosts in vCenter and then adding the entire cluster to CloudStack. See [Add Cluster: vSphere](#) .
 - (Versão 3.0.0 só, em 3.0.1, este campo é somente leitura) Escolha o tipo de software hypervisor que todos os anfitriões deste cluster será executado. Se você escolher a VMware, campos adicionais aparecem assim que você pode dar informações sobre um cluster vSphere. Para os servidores vSphere, recomendamos criar o cluster de hosts em vCenter e adicionando todo o cluster para CloudStack. Consulte [Adicionar Cluster: vSphere](#).
9. In a new cluster, CloudStack adds the first host for you. You can always add more hosts later. For an overview of what a host is, see [Seção 2.5, “Sobre hosts”](#).



Nota

When you deploy CloudStack, the hypervisor host must not have any VMs already running.

Antes de configurar o host, você precisa instalar o software hipervisor no host. Você precisará saber qual versão do software hipervisor é suportada pelo CloudStack e qual configuração adicional é requerida para

- Citrix XenServer Installation for CloudStack
- Instalação e configuração do VMware vSphere
- KVM Installation and Configuration

Para configurar o primeiro host, Entre o seguinte, então clique Next:

- **Host Name.** O nome DNS ou endereço IP do host..
- **Username.** Usually root.
- **Password.** Esta é a senha para o usuário chamado acima (do seu XenServer ou instalação KVM).
- **Host Tags.** (Optional) Any labels that you use to categorize hosts for ease of maintenance. For example, you can set to the cloud's HA tag (set in the ha.tag global configuration parameter) if you want this host to be used only for VMs with the "high availability" feature enabled. For more information, see HA-Enabled Virtual Machines as well as HA for Hosts, both in the Administration Guide.

10. In a new cluster, CloudStack adds the first primary storage server for you. You can always add more servers later. For an overview of what primary storage is, see [Seção 2.6, "Sobre storage primária"](#).

Para configurar o primeiro servidor de storage primário, entre o seguinte, então clique Next.

- **Name.** O nome do dispositivo do storage.
- **Protocol.** For XenServer, choose either NFS, iSCSI, or PreSetup. For KVM, choose NFS, SharedMountPoint, CLVM, and RBD. For vSphere choose either VMFS (iSCSI or FiberChannel) or NFS. The remaining fields in the screen vary depending on what you choose here.

NFS	<ul style="list-style-type: none">• Server. The IP address or DNS name of the storage device.• Path. The exported path from the server.• Tags (optional). The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. <p>The tag sets on primary storage across clusters in a Zone must be identical. For</p>
-----	---

	<p>example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.</p>
iSCSI	<ul style="list-style-type: none"> • Server. The IP address or DNS name of the storage device. • Target IQN. The IQN of the target. For example, iqn.1986-03.com.sun:02:01ec9bb549-1271378984. • Lun. The LUN number. For example, 3. • Tags (optional). The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. <p>The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.</p>
preSetup	<ul style="list-style-type: none"> • Server. The IP address or DNS name of the storage device. • SR Name-Label. Enter the name-label of the SR that has been set up outside CloudStack. • Tags (optional). The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. <p>The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.</p>
SharedMountPoint	<ul style="list-style-type: none"> • Path. The path on each host that is where this primary storage is mounted. For example, "/mnt/primary". • Tags (optional). The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. <p>The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in</p>

	the Zone must also provide primary storage that has tags T1 and T2.
VMFS	<ul style="list-style-type: none"> • Server. The IP address or DNS name of the vCenter server. • Path. A combination of the datacenter name and the datastore name. The format is "/" datacenter name "/" datastore name. For example, "/cloud.dc.VM/cluster1datastore". • Tags (optional). The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings. <p>The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.</p>

11. In a new zone, CloudStack adds the first secondary storage server for you. For an overview of what secondary storage is, see [Seção 2.7, “Sobre storage secundária”](#).

Before you can fill out this screen, you need to prepare the secondary storage by setting up NFS shares and installing the latest CloudStack System VM template. See Adding Secondary Storage :

- **NFS Server.** The IP address of the server or fully qualified domain name of the server.
- **Path.** The exported path from the server.

12. Click Launch.

7.4. Adicionando um pod

Quando você cria uma nova zona, o CloudStack adiciona o primeiro pod para você. Você pode adicionar mais pods a qualquer momento usando o procedimento nesta seção.

1. Faça login na interface de usuário do CloudStack. Veja [Seção 5.1, “Login na interface de usuário”](#).
2. No painel de navegação à esquerda, selecione Infrastructure. Em Zones, clique View More, então clique na zona à qual você deseja adicionar um pod.
3. Clique na aba Compute and Storage. No nó Pods do diagrama, clique em View All.
4. Clique em Add Pod.
5. Forneça os seguintes detalhes no diálogo.
 - **Name.** O nome do pod.
 - **Gateway.** O gateway para os hosts no pod.
 - **Netmask.** O prefixo de rede que define a subrede deste pod. Utilize a notação CIDR.

- **Start/End Reserved System IP.** O intervalo IP na rede de gerenciamento que o CloudStack usa para gerenciar várias máquinas virtuais de sistema, tais como as máquinas virtuais de storage secundária, máquinas virtuais de proxy de console, e DHCP. Para mais informações, veja Endereços IP reservados pelo sistema.

6. Clique em OK.

7.5. Adicionando um cluster

Você precisa informar ao CloudStack sobre os hosts que ele irá gerenciar. Hosts existem em clusters, portanto antes de você começar a adicionar hosts à nuvem, você deve adicionar pelo menos um cluster.

7.5.1. Add Cluster: KVM or XenServer

These steps assume you have already installed the hypervisor on the hosts and logged in to the CloudStack UI.

1. In the left navigation, choose Infrastructure. In Zones, click View More, then click the zone in which you want to add the cluster.
2. Click the Compute tab.
3. In the Clusters node of the diagram, click View All.
4. Click Add Cluster.
5. Choose the hypervisor type for this cluster.
6. Choose the pod in which you want to create the cluster.
7. Enter a name for the cluster. This can be text of your choosing and is not used by CloudStack.
8. Clique em OK.

7.5.2. Add Cluster: vSphere


Host management for vSphere is done through a combination of vCenter and the CloudStack admin UI. CloudStack requires that all hosts be in a CloudStack cluster, but the cluster may consist of a single host. As an administrator you must decide if you would like to use clusters of one host or of multiple hosts. Clusters of multiple hosts allow for features like live migration. Clusters also require shared storage such as NFS or iSCSI.

For vSphere servers, we recommend creating the cluster of hosts in vCenter and then adding the entire cluster to CloudStack. Follow these requirements:

- Do not put more than 8 hosts in a vSphere cluster
- Make sure the hypervisor hosts do not have any VMs already running before you add them to CloudStack.

To add a vSphere cluster to CloudStack:

1. Create the cluster of hosts in vCenter. Follow the vCenter instructions to do this. You will create a cluster that looks something like this in vCenter.



2. Log in to the UI.
3. In the left navigation, choose Infrastructure. In Zones, click View More, then click the zone in which you want to add the cluster.
4. Click the Compute tab, and click View All on Pods. Choose the pod to which you want to add the cluster.
5. Click View Clusters.
6. Click Add Cluster.
7. In Hypervisor, choose VMware.
8. Provide the following information in the dialog. The fields below make reference to values from vCenter.
 - Cluster Name. Enter the name of the cluster you created in vCenter. For example, "cloud.cluster.2.2.1"
 - vCenter Host. Enter the hostname or IP address of the vCenter server.
 - vCenter Username. Enter the username that CloudStack should use to connect to vCenter. This user must have all administrative privileges.
 - vCenter Password. Enter the password for the user named above
 - vCenter Datacenter. Enter the vCenter datacenter that the cluster is in. For example, "cloud.dc.VM".

The screenshot shows a dialog box titled "Add Cluster" with the following fields and values:

- Zone: ZONE-NEXUS-ADV
- Hypervisor: VMware
- Pod: POD-1
- Cluster Name: doc-cluster
- vCenter Host: host-cs-vcenter
- vCenter Username: admin
- vCenter Password: [masked]
- vCenter Datacenter: doc-datacenter
- Nexus dvSwitch IP Address: 10.10.105.10
- Nexus dvSwitch Username: dv-admin
- Nexus dvSwitch Password: [masked]

Buttons: Cancel, OK

There might be a slight delay while the cluster is provisioned. It will automatically display in the UI

7.6. Adding a Host

1. Before adding a host to the CloudStack configuration, you must first install your chosen hypervisor on the host. CloudStack can manage hosts running VMs under a variety of hypervisors.

The CloudStack Installation Guide provides instructions on how to install each supported hypervisor and configure it for use with CloudStack. See the appropriate section in the Installation Guide for information about which version of your chosen hypervisor is supported, as well as crucial additional steps to configure the hypervisor hosts for use with CloudStack.



Atenção

Be sure you have performed the additional CloudStack-specific configuration steps described in the hypervisor installation section for your particular hypervisor.

2. Now add the hypervisor host to CloudStack. The technique to use varies depending on the hypervisor.
 - [Seção 7.6.1, “Adding a Host \(XenServer or KVM\)”](#)
 - [Seção 7.6.2, “Adding a Host \(vSphere\)”](#)

7.6.1. Adding a Host (XenServer or KVM)

XenServer and KVM hosts can be added to a cluster at any time.

7.6.1.1. Requirements for XenServer and KVM Hosts



Atenção

Make sure the hypervisor host does not have any VMs already running before you add it to CloudStack.

Configuration requirements:

- Each cluster must contain only hosts with the identical hypervisor.
- For XenServer, do not put more than 8 hosts in a cluster.
- For KVM, do not put more than 16 hosts in a cluster.

For hardware requirements, see the installation section for your hypervisor in the CloudStack Installation Guide.

7.6.1.1.1. XenServer Host Additional Requirements

If network bonding is in use, the administrator must cable the new host identically to other hosts in the cluster.

For all additional hosts to be added to the cluster, run the following command. This will cause the host to join the master in a XenServer pool.

```
# xe pool-join master-address=[master IP] master-username=root master-password=[your password]
```



Nota

Ao copiar e colar um comando, certifique-se que o comando tenha colado como uma única linha antes de executar. Alguns viewers documento pode introduzir quebras de linha indesejadas no texto copiado.

With all hosts added to the XenServer pool, run the cloud-setup-bond script. This script will complete the configuration and setup of the bonds on the new hosts in the cluster.

1. Copy the script from the Management Server in `/usr/lib64/cloud/common/scripts/vm/hypervisor/xenserver/cloud-setup-bonding.sh` to the master host and ensure it is executable.
2. Run the script:

```
# ./cloud-setup-bonding.sh
```

7.6.1.1.2. KVM Host Additional Requirements

- If shared mountpoint storage is in use, the administrator should ensure that the new host has all the same mountpoints (with storage mounted) as the other hosts in the cluster.
- Make sure the new host has the same network configuration (guest, private, and public network) as other hosts in the cluster.
- If you are using OpenVswitch bridges edit the file `agent.properties` on the KVM host and set the parameter `network.bridge.type` to `openvswitch` before adding the host to CloudStack

7.6.1.2. Adding a XenServer or KVM Host

- If you have not already done so, install the hypervisor software on the host. You will need to know which version of the hypervisor software version is supported by CloudStack and what additional configuration is required to ensure the host will work with CloudStack. To find these installation details, see the appropriate section for your hypervisor in the CloudStack Installation Guide.
- Faça login na interface de usuário do CloudStack como administrador.
- In the left navigation, choose Infrastructure. In Zones, click View More, then click the zone in which you want to add the host.
- Click the Compute tab. In the Clusters node, click View All.
- Click the cluster where you want to add the host.
- Click View Hosts.
- Click Add Host.
- Provide the following information.
 - Host Name. The DNS name or IP address of the host.
 - Username. Usually root.
 - Password. This is the password for the user from your XenServer or KVM install).
 - Host Tags (Optional). Any labels that you use to categorize hosts for ease of maintenance. For example, you can set to the cloud's HA tag (set in the `ha.tag` global configuration parameter) if you want this host to be used only for VMs with the "high availability" feature enabled. For more information, see HA-Enabled Virtual Machines as well as HA for Hosts.

There may be a slight delay while the host is provisioned. It should automatically display in the UI.

- Repeat for additional hosts.

7.6.2. Adding a Host (vSphere)

For vSphere servers, we recommend creating the cluster of hosts in vCenter and then adding the entire cluster to CloudStack. See Add Cluster: vSphere.

7.7. Adicionar Storage Primário

7.7.1. System Requirements for Primary Storage

Hardware requirements:

- Any standards-compliant iSCSI or NFS server that is supported by the underlying hypervisor.
- The storage server should be a machine with a large number of disks. The disks should ideally be managed by a hardware RAID controller.
- Minimum required capacity depends on your needs.

When setting up primary storage, follow these restrictions:

- Primary storage cannot be added until a host has been added to the cluster.
- If you do not provision shared primary storage, you must set the global configuration parameter `system.vm.local.storage.required` to true, or else you will not be able to start VMs.

7.7.2. Adding Primary Storage

When you create a new zone, the first primary storage is added as part of that procedure. You can add primary storage servers at any time, such as when adding a new cluster or adding more servers to an existing cluster.



Atenção

Be sure there is nothing stored on the server. Adding the server to CloudStack will destroy any existing data.

1. Log in to the CloudStack UI (see [Seção 5.1, “Login na interface de usuário”](#)).
2. In the left navigation, choose Infrastructure. In Zones, click View More, then click the zone in which you want to add the primary storage.
3. Click the Compute tab.
4. In the Primary Storage node of the diagram, click View All.
5. Click Add Primary Storage.
6. Provide the following information in the dialog. The information required varies depending on your choice in Protocol.
 - **Pod.** The pod for the storage device.
 - **Cluster.** The cluster for the storage device.
 - **Name.** O nome do dispositivo do storage.
 - **Protocol.** For XenServer, choose either NFS, iSCSI, or PreSetup. For KVM, choose NFS or SharedMountPoint. For vSphere choose either VMFS (iSCSI or FiberChannel) or NFS.

- **Server (for NFS, iSCSI, or PreSetup).** The IP address or DNS name of the storage device.
- **Server (for VMFS).** The IP address or DNS name of the vCenter server.
- **Path (for NFS).** In NFS this is the exported path from the server.
- **Path (for VMFS).** In vSphere this is a combination of the datacenter name and the datastore name. The format is "/" datacenter name "/" datastore name. For example, "/cloud.dc.VM/cluster1datastore".
- **Path (for SharedMountPoint).** With KVM this is the path on each host that is where this primary storage is mounted. For example, "/mnt/primary".
- **SR Name-Label (for PreSetup).** Enter the name-label of the SR that has been set up outside CloudStack.
- **Target IQN (for iSCSI).** In iSCSI this is the IQN of the target. For example, `iqn.1986-03.com.sun:02:01ec9bb549-1271378984`.
- **Lun # (for iSCSI).** In iSCSI this is the LUN number. For example, 3.
- **Tags (optional).** The comma-separated list of tags for this storage device. It should be an equivalent set or superset of the tags on your disk offerings..

The tag sets on primary storage across clusters in a Zone must be identical. For example, if cluster A provides primary storage that has tags T1 and T2, all other clusters in the Zone must also provide primary storage that has tags T1 and T2.

7. Clique em OK.

7.8. Adicionar Storage Secundário

7.8.1. System Requirements for Secondary Storage

- NFS storage appliance or Linux NFS server
- (Optional) OpenStack Object Storage (Swift) (see <http://swift.openstack.org>)
- 100GB minimum capacity
- A secondary storage device must be located in the same zone as the guest VMs it serves.
- Each Secondary Storage server must be available to all hosts in the zone.

7.8.2. Adding Secondary Storage

When you create a new zone, the first secondary storage is added as part of that procedure. You can add secondary storage servers at any time to add more servers to an existing zone.



Atenção

Be sure there is nothing stored on the server. Adding the server to CloudStack will destroy any existing data.

1. If you are going to use Swift for cloud-wide secondary storage, you must add the Swift storage to CloudStack before you add the local zone secondary storage servers. See [Seção 7.3, “Adicionando uma zona”](#).
2. To prepare for local zone secondary storage, you should have created and mounted an NFS share during Management Server installation. See Preparing NFS Shares in the Installation Guide.
3. Make sure you prepared the system VM template during Management Server installation. See Prepare the System VM Template in the Installation Guide.
4. Now that the secondary storage server for per-zone storage is prepared, add it to CloudStack. Secondary storage is added as part of the procedure for adding a new zone. See [Seção 7.3, “Adicionando uma zona”](#).

7.9. Initialize and Test

After everything is configured, CloudStack will perform its initialization. This can take 30 minutes or more, depending on the speed of your network. When the initialization has completed successfully, the administrator's Dashboard should be displayed in the CloudStack UI.

1. Verify that the system is ready. In the left navigation bar, select Templates. Click on the CentOS 5.5 (64bit) no Gui (KVM) template. Check to be sure that the status is "Download Complete." Do not proceed to the next step until this status is displayed.
2. Go to the Instances tab, and filter by My Instances.
3. Click Add Instance and follow the steps in the wizard.
 - a. Choose the zone you just added.
 - b. In the template selection, choose the template to use in the VM. If this is a fresh installation, likely only the provided CentOS template is available.
 - c. Select a service offering. Be sure that the hardware you have allows starting the selected service offering.
 - d. In data disk offering, if desired, add another data disk. This is a second volume that will be available to but not mounted in the guest. For example, in Linux on XenServer you will see /dev/xvdb in the guest after rebooting the VM. A reboot is not required if you have a PV-enabled OS kernel in use.
 - e. In default network, choose the primary network for the guest. In a trial installation, you would have only one option here.
 - f. Optionally give your VM a name and a group. Use any descriptive text you would like.

-
- g. Click Launch VM. Your VM will be created and started. It might take some time to download the template and complete the VM startup. You can watch the VM's progress in the Instances screen.

- 4.

To use the VM, click the View Console button.



For more information about using VMs, including instructions for how to allow incoming network traffic to the VM, start, stop, and delete VMs, and move a VM from one host to another, see [Working With Virtual Machines in the Administrator's Guide](#).

Congratulations! You have successfully completed a CloudStack Installation.

If you decide to grow your deployment, you can add more hosts, primary storage, zones, pods, and clusters.

Ofertas de serviços

Neste capítulo discutimos ofertas serviços de computação, disco, e sistema. Ofertas de rede são discutidas na seção sobre configuração de rede para usuários.

8.1. Compute and Disk Service Offerings

A service offering is a set of virtual hardware features such as CPU core count and speed, memory, and disk size. The CloudStack administrator can set up various offerings, and then end users choose from the available offerings when they create a new VM. A service offering includes the following elements:

- CPU, memory, and network resource guarantees
- How resources are metered
- How the resource usage is charged
- How often the charges are generated

For example, one service offering might allow users to create a virtual machine instance that is equivalent to a 1 GHz Intel® Core™ 2 CPU, with 1 GB memory at \$0.20/hour, with network traffic metered at \$0.10/GB. Based on the user's selected offering, CloudStack emits usage records that can be integrated with billing systems. CloudStack separates service offerings into compute offerings and disk offerings. The computing service offering specifies:

- Guest CPU
- Guest RAM
- Guest Networking type (virtual or direct)
- Tags on the root disk

The disk offering specifies:

- Disk size (optional). An offering without a disk size will allow users to pick their own
- Tags on the data disk

8.1.1. Creating a New Compute Offering

To create a new compute offering:

1. Log in with admin privileges to the CloudStack UI.
2. In the left navigation bar, click Service Offerings.
3. In Select Offering, choose Compute Offering.
4. Click Add Compute Offering.
5. In the dialog, make the following choices:
 - **Name:** Any desired name for the service offering.
 - **Description:** A short description of the offering that can be displayed to users

- **Storage type:** The type of disk that should be allocated. Local allocates from storage attached directly to the host where the system VM is running. Shared allocates from storage accessible via NFS.
- **# of CPU cores:** The number of cores which should be allocated to a system VM with this offering
- **CPU (in MHz):** The CPU speed of the cores that the system VM is allocated. For example, "2000" would provide for a 2 GHz clock.
- **Memory (in MB):** The amount of memory in megabytes that the system VM should be allocated. For example, "2048" would provide for a 2 GB RAM allocation.
- **Network Rate:** Allowed data transfer rate in MB per second.
- **Offer HA:** If yes, the administrator can choose to have the system VM be monitored and as highly available as possible.
- **Storage Tags:** The tags that should be associated with the primary storage used by the system VM.
- **Host Tags:** (Optional) Any tags that you use to organize your hosts
- **CPU cap:** Whether to limit the level of CPU usage even if spare capacity is available.
- **Public:** Indicate whether the service offering should be available all domains or only some domains. Choose Yes to make it available to all domains. Choose No to limit the scope to a subdomain; CloudStack will then prompt for the subdomain's name.

6. Clique em Add.

8.1.2. Creating a New Disk Offering

To create a system service offering:

1. Log in with admin privileges to the CloudStack UI.
2. In the left navigation bar, click Service Offerings.
3. In Select Offering, choose Disk Offering.
4. Click Add Disk Offering.
5. In the dialog, make the following choices:
 - Name. Any desired name for the system offering.
 - Description. A short description of the offering that can be displayed to users
 - Custom Disk Size. If checked, the user can set their own disk size. If not checked, the root administrator must define a value in Disk Size.
 - Disk Size. Appears only if Custom Disk Size is not selected. Define the volume size in GB.
 - (Optional)Storage Tags. The tags that should be associated with the primary storage for this disk. Tags are a comma separated list of attributes of the storage. For example "ssd,blue". Tags are also added on Primary Storage. CloudStack matches tags on a disk offering to tags on the storage. If a tag is present on a disk offering that tag (or tags) must also be present on Primary

Storage for the volume to be provisioned. If no such primary storage exists, allocation from the disk offering will fail..

- **Public.** Indicate whether the service offering should be available all domains or only some domains. Choose Yes to make it available to all domains. Choose No to limit the scope to a subdomain; CloudStack will then prompt for the subdomain's name.

6. Clique em Add.

8.1.3. Modifying or Deleting a Service Offering

Service offerings cannot be changed once created. This applies to both compute offerings and disk offerings.

A service offering can be deleted. If it is no longer in use, it is deleted immediately and permanently. If the service offering is still in use, it will remain in the database until all the virtual machines referencing it have been deleted. After deletion by the administrator, a service offering will not be available to end users that are creating new instances.

8.2. System Service Offerings

System service offerings provide a choice of CPU speed, number of CPUs, tags, and RAM size, just as other service offerings do. But rather than being used for virtual machine instances and exposed to users, system service offerings are used to change the default properties of virtual routers, console proxies, and other system VMs. System service offerings are visible only to the CloudStack root administrator. CloudStack provides default system service offerings. The CloudStack root administrator can create additional custom system service offerings.

When CloudStack creates a virtual router for a guest network, it uses default settings which are defined in the system service offering associated with the network offering. You can upgrade the capabilities of the virtual router by applying a new network offering that contains a different system service offering. All virtual routers in that network will begin using the settings from the new service offering.

8.2.1. Creating a New System Service Offering

To create a system service offering:

1. Log in with admin privileges to the CloudStack UI.
2. In the left navigation bar, click Service Offerings.
3. In Select Offering, choose System Offering.
4. Click Add System Service Offering.
5. In the dialog, make the following choices:
 - **Name.** Any desired name for the system offering.
 - **Description.** A short description of the offering that can be displayed to users
 - **System VM Type.** Select the type of system virtual machine that this offering is intended to support.

- Storage type. The type of disk that should be allocated. Local allocates from storage attached directly to the host where the system VM is running. Shared allocates from storage accessible via NFS.
- # of CPU cores. The number of cores which should be allocated to a system VM with this offering
- CPU (in MHz). The CPU speed of the cores that the system VM is allocated. For example, "2000" would provide for a 2 GHz clock.
- Memory (in MB). The amount of memory in megabytes that the system VM should be allocated. For example, "2048" would provide for a 2 GB RAM allocation.
- Network Rate. Allowed data transfer rate in MB per second.
- Offer HA. If yes, the administrator can choose to have the system VM be monitored and as highly available as possible.
- Storage Tags. The tags that should be associated with the primary storage used by the system VM.
- Host Tags. (Optional) Any tags that you use to organize your hosts
- CPU cap. Whether to limit the level of CPU usage even if spare capacity is available.
- Public. Indicate whether the service offering should be available all domains or only some domains. Choose Yes to make it available to all domains. Choose No to limit the scope to a subdomain; CloudStack will then prompt for the subdomain's name.

6. Clique em Add.

8.3. Network Throttling

Network throttling is the process of controlling the network access and bandwidth usage based on certain rules. CloudStack controls this behaviour of the guest networks in the cloud by using the network rate parameter. This parameter is defined as the default data transfer rate in Mbps (Megabits Per Second) allowed in a guest network. It defines the upper limits for network utilization. If the current utilization is below the allowed upper limits, access is granted, else revoked.

You can throttle the network bandwidth either to control the usage above a certain limit for some accounts, or to control network congestion in a large cloud environment. The network rate for your cloud can be configured on the following:

- Network Offering
- Plano
- Global parameter

If network rate is set to NULL in service offering, the value provided in the `vm.network.throttling.rate` global parameter is applied. If the value is set to NULL for network offering, the value provided in the `network.throttling.rate` global parameter is considered.

For the default public, storage, and management networks, network rate is set to 0. This implies that the public, storage, and management networks will have unlimited bandwidth by default. For default guest networks, network rate is set to NULL. In this case, network rate is defaulted to the global parameter value.

The following table gives you an overview of how network rate is applied on different types of networks in CloudStack.

Redes	Network Rate Is Taken from
Guest network of Virtual Router	Guest Network Offering
Public network of Virtual Router	Guest Network Offering
Storage network of Secondary Storage VM	System Network Offering
Management network of Secondary Storage VM	System Network Offering
Storage network of Console Proxy VM	System Network Offering
Management network of Console Proxy VM	System Network Offering
Storage network of Virtual Router	System Network Offering
Management network of Virtual Router	System Network Offering
Public network of Secondary Storage VM	System Network Offering
Public network of Console Proxy VM	System Network Offering
Default network of a guest VM	Compute Offering
Additional networks of a guest VM	Corresponding Network Offerings

A guest VM must have a default network, and can also have many additional networks. Depending on various parameters, such as the host and virtual switch used, you can observe a difference in the network rate in your cloud. For example, on a VMware host the actual network rate varies based on where they are configured (compute offering, network offering, or both); the network type (shared or isolated); and traffic direction (ingress or egress).

The network rate set for a network offering used by a particular network in CloudStack is used for the traffic shaping policy of a port group, for example: port group A, for that network: a particular subnet or VLAN on the actual network. The virtual routers for that network connects to the port group A, and by default instances in that network connects to this port group. However, if an instance is deployed with a compute offering with the network rate set, and if this rate is used for the traffic shaping policy of another port group for the network, for example port group B, then instances using this compute offering are connected to the port group B, instead of connecting to port group A.

The traffic shaping policy on standard port groups in VMware only applies to the egress traffic, and the net effect depends on the type of network used in CloudStack. In shared networks, ingress traffic is unlimited for CloudStack, and egress traffic is limited to the rate that applies to the port group used by the instance if any. If the compute offering has a network rate configured, this rate applies to the egress traffic, otherwise the network rate set for the network offering applies. For isolated networks, the network rate set for the network offering, if any, effectively applies to the ingress traffic. This is mainly because the network rate set for the network offering applies to the egress traffic from the virtual router to the instance. The egress traffic is limited by the rate that applies to the port group used by the instance if any, similar to shared networks.

For example:

Network rate of network offering = 10 Mbps

Network rate of compute offering = 200 Mbps

In shared networks, ingress traffic will not be limited for CloudStack, while egress traffic will be limited to 200 Mbps. In an isolated network, ingress traffic will be limited to 10 Mbps and egress to 200 Mbps.

8.4. Changing the Default System Offering for System VMs

You can manually change the system offering for a particular System VM. Additionally, as a CloudStack administrator, you can also change the default system offering used for System VMs.

1. Create a new system offering.

For more information, see [Creating a New System Service Offering](#).

2. Back up the database:

```
mysqldump -u root -p cloud | bzip2 > cloud_backup.sql.bz2
```

3. Open an MySQL prompt:

```
mysql -u cloud -p cloud
```

4. Run the following queries on the cloud database.

- a. In the `disk_offering` table, identify the original default offering and the new offering you want to use by default.

Take a note of the ID of the new offering.

```
select id,name,unique_name,type from disk_offering;
```

- b. For the original default offering, set the value of `unique_name` to `NULL`.

```
# update disk_offering set unique_name = NULL where id = 10;
```

Ensure that you use the correct value for the ID.

- c. For the new offering that you want to use by default, set the value of `unique_name` as follows:

For the default Console Proxy VM (CPVM) offering, set `unique_name` to 'Cloud.com-ConsoleProxy'. For the default Secondary Storage VM (SSVM) offering, set `unique_name` to 'Cloud.com-SecondaryStorage'. For example:

```
update disk_offering set unique_name = 'Cloud.com-ConsoleProxy' where id = 16;
```

5. Restart CloudStack Management Server. Restarting is required because the default offerings are loaded into the memory at startup.

```
service cloudstack-management restart
```

6. Destroy the existing CPVM or SSVM offerings and wait for them to be recreated. The new CPVM or SSVM are configured with the new offering.

Setting Up Networking for Users

9.1. Overview of Setting Up Networking for Users

People using cloud infrastructure have a variety of needs and preferences when it comes to the networking services provided by the cloud. As a CloudStack administrator, you can do the following things to set up networking for your users:

- Set up physical networks in zones
- Set up several different providers for the same service on a single physical network (for example, both Cisco and Juniper firewalls)
- Bundle different types of network services into network offerings, so users can choose the desired network services for any given virtual machine
- Add new network offerings as time goes on so end users can upgrade to a better class of service on their network
- Provide more ways for a network to be accessed by a user, such as through a project of which the user is a member

9.2. About Virtual Networks

A virtual network is a logical construct that enables multi-tenancy on a single physical network. In CloudStack a virtual network can be shared or isolated.

9.2.1. Isolated Networks

An isolated network can be accessed only by virtual machines of a single account. Isolated networks have the following properties.

- Resources such as VLAN are allocated and garbage collected dynamically
- There is one network offering for the entire network
- The network offering can be upgraded or downgraded but it is for the entire network

9.2.2. Shared Networks

A shared network can be accessed by virtual machines that belong to many different accounts. Network Isolation on shared networks is accomplished using techniques such as security groups (supported only in basic zones).

- Shared Networks are created by the administrator
- Shared Networks can be designated to a certain domain
- Shared Network resources such as VLAN and physical network that it maps to are designated by the administrator
- Shared Networks are isolated by security groups
- Public Network is a shared network that is not shown to the end users

9.2.3. Runtime Allocation of Virtual Network Resources

When you define a new virtual network, all your settings for that network are stored in CloudStack. The actual network resources are activated only when the first virtual machine starts in the network. When all virtual machines have left the virtual network, the network resources are garbage collected so they can be allocated again. This helps to conserve network resources.

9.3. Provedores de Serviços de Rede



Nota

For the most up-to-date list of supported network service providers, see the CloudStack UI or call `listNetworkServiceProviders`.

A service provider (also called a network element) is hardware or virtual appliance that makes a network service possible; for example, a firewall appliance can be installed in the cloud to provide firewall service. On a single network, multiple providers can provide the same network service. For example, a firewall service may be provided by Cisco or Juniper devices in the same physical network.

You can have multiple instances of the same service provider in a network (say, more than one Juniper SRX device).

If different providers are set up to provide the same service on the network, the administrator can create network offerings so users can specify which network service provider they prefer (along with the other choices offered in network offerings). Otherwise, CloudStack will choose which provider to use whenever the service is called for.

Supported Network Service Providers

CloudStack ships with an internal list of the supported service providers, and you can choose from this list when creating a network offering.

	Roteador Virtual	Citrix NetScaler	Juniper SRX	F5 BigIP	Host based (KVM/Xen)
Remote Access VPN	Sim	Não	Não	Não	Não
DNS/DHCP/ User Data	Sim	Não	Não	Não	Não
Firewall	Sim	Não	Sim	Não	Não
Balanceamento de Carga	Sim	Sim	Não	Sim	Não
IP Elástico	Não	Sim	Não	Não	Não
LB Elástico	Não	Sim	Não	Não	Não
NAT de origem	Sim	Não	Sim	Não	Não
NAT estática	Sim	Sim	Sim	Não	Não
Encaminhamento de Porta	Sim	Não	Sim	Não	Não

9.4. Oferta de Rede



Nota

Para a mais atualizada lista de serviços de rede suportados, veja o interface de usuário do CloudStack ou chame `listNetworkServices`.

Uma oferta de rede é um conjunto identificado de serviços de rede, tais como:

- DHCP
- DNS
- Source NAT
- NAT Estático
- Encaminhamento de Porta
- Balanceamento de Carga
- Firewall
- VPN
- Opcional) Nome de um entre vários fornecedores para usar para um dado serviço, tal como Juniper para o firewall
- (Opcional) Tag de rede para especificar qual a rede física a utilizar

Ao criar uma nova máquina virtual, o usuário escolhe uma das ofertas de rede disponíveis, e isto determina quais serviços de rede a máquina virtual pode usar.

O administrador do CloudStack pode criar qualquer número de ofertas de rede customizadas, em adição às ofertas de rede default providas pelo CloudStack. Criando múltiplas ofertas de rede customizadas, você pode configurar sua nuvem para oferecer diferentes classes de serviço em uma única rede física multi-tenant. Por exemplo, enquanto o cabeamento físico subjacente pode ser o mesmo para dois tenants, o tenant A pode requerer apenas proteção de firewall simples para o seu website, enquanto o tenant B pode estar executando um web server farm e requerer uma solução escalável de firewall, solução de balanceamento de carga e redes alternativas para acessar o database no backend.



Nota

Se você cria regras de balanceamento de carga enquanto usando um oferta de serviço de rede que inclui um equipamento externo de balanceamento de carga, como o NetScaler, e depois altera a oferta de serviço para um que usa o roteador virtual do CloudStack, você deve criar uma regra no firewall do roteador virtual para cada uma das regras de balanceamento de carga existentes, de forma que elas possam continuar funcionando.

Ao criar uma nova rede virtual, o administrador do CloudStack escolhe qual oferta de rede habilitar para aquela rede. Cada rede virtual é associada com uma oferta de rede. Uma rede virtual pode ser incrementada ou decrementada alterando sua oferta de rede associada. Se você fizer isto, certifique-se de reprogramar a rede física para coincidir.

O CloudStack também possui ofertas de rede internas para uso das máquinas virtuais de sistema do CloudStack. Estas ofertas de rede não são visíveis pelos usuários, mas podem ser modificadas pelos administradores.

9.4.1. Creating a New Network Offering

To create a network offering:

1. Log in with admin privileges to the CloudStack UI.
2. In the left navigation bar, click Service Offerings.
3. In Select Offering, choose Network Offering.
4. Click Add Network Offering.
5. In the dialog, make the following choices:
 - **Name.** Any desired name for the network offering.
 - **Description.** A short description of the offering that can be displayed to users.
 - **Network Rate.** Allowed data transfer rate in MB per second.
 - **Guest Type.** Choose whether the guest network is isolated or shared.

For a description of this term, see [Seção 9.2, “About Virtual Networks”](#).

- **Persistent.** Indicate whether the guest network is persistent or not. The network that you can provision without having to deploy a VM on it is termed persistent network. For more information, see [Seção 15.20, “Persistent Networks”](#).
- **Specify VLAN.** (Isolated guest networks only) Indicate whether a VLAN should be specified when this offering is used.
- **VPC.** This option indicate whether the guest network is Virtual Private Cloud-enabled. A Virtual Private Cloud (VPC) is a private, isolated part of CloudStack. A VPC can have its own virtual network topology that resembles a traditional physical network. For more information on VPCs, see [Seção 15.19.1, “About Virtual Private Clouds”](#).
- **Supported Services.** Select one or more of the possible network services. For some services, you must also choose the service provider; for example, if you select Load Balancer, you can choose the CloudStack virtual router or any other load balancers that have been configured in the cloud. Depending on which services you choose, additional fields may appear in the rest of the dialog box.

Based on the guest network type selected, you can see the following supported services:


Serviços Suportados	Descrição	Isolated	Compatilhado
DHCP	For more information, see Seção 15.16 , “DNS e DHCP”.	Supported	Supported
DNS	For more information, see Seção 15.16 , “DNS e DHCP”.	Supported	Supported
Load Balancer	If you select Load Balancer, you can choose the CloudStack virtual router or any other load balancers that have been configured in the cloud.	Supported	Supported
Firewall	For more information, see the Administration Guide.	Supported	Supported
Source NAT	If you select Source NAT, you can choose the CloudStack virtual router or any other Source NAT providers that have been configured in the cloud.	Supported	Supported
NAT Estático	If you select Static NAT, you can choose the CloudStack virtual router or any other Static NAT providers that have been configured in the cloud.	Supported	Supported
Encaminhamento de Porta	If you select Port Forwarding, you can choose the CloudStack virtual router or any other Port Forwarding providers that have been configured in the cloud.	Supported	Not Supported
VPN	For more information, see Seção 15.17 , “VPN”.	Supported	Not Supported
User Data	For more information, see Seção 20.3 ,	Not Supported	Supported

Serviços Suportados	Descrição	Isolated	Compatilhado
	<i>“User Data and Meta Data”.</i>		
Network ACL	For more information, see <i>Seção 15.19.4, “Configuring Access Control List”.</i>	Supported	Not Supported
Security Groups	For more information, see <i>Seção 15.7.2, “Adicionando um grupo de segurança”.</i>	Not Supported	Supported

- **System Offering.** If the service provider for any of the services selected in Supported Services is a virtual router, the System Offering field appears. Choose the system service offering that you want virtual routers to use in this network. For example, if you selected Load Balancer in Supported Services and selected a virtual router to provide load balancing, the System Offering field appears so you can choose between the CloudStack default system service offering and any custom system service offerings that have been defined by the CloudStack root administrator.

For more information, see *Seção 8.2, “System Service Offerings”.*

- **Redundant router capability.** Available only when Virtual Router is selected as the Source NAT provider. Select this option if you want to use two virtual routers in the network for uninterrupted connection: one operating as the master virtual router and the other as the backup. The master virtual router receives requests from and sends responses to the user’s VM. The backup virtual router is activated only when the master is down. After the failover, the backup becomes the master virtual router. CloudStack deploys the routers on different hosts to ensure reliability if one host is down.
- **Conserve mode.** Indicate whether to use conserve mode. In this mode, network resources are allocated only when the first virtual machine starts in the network. When conservative mode is off, the public IP can only be used for a single service. For example, a public IP used for a port forwarding rule cannot be used for defining other services, such as StaticNAT or load balancing. When the conserve mode is on, you can define more than one service on the same public IP.



Nota

If StaticNAT is enabled, irrespective of the status of the conserve mode, no port forwarding or load balancing rule can be created for the IP. However, you can add the firewall rules by using the createFirewallRule command.

- **Tags.** Network tag to specify which physical network to use.

6. Clique em Add.

Working With Virtual Machines

10.1. About Working with Virtual Machines

CloudStack provides administrators with complete control over the lifecycle of all guest VMs executing in the cloud. CloudStack provides several guest management operations for end users and administrators. VMs may be stopped, started, rebooted, and destroyed.

Guest VMs have a name and group. VM names and groups are opaque to CloudStack and are available for end users to organize their VMs. Each VM can have three names for use in different contexts. Only two of these names can be controlled by the user:

- Instance name – a unique, immutable ID that is generated by CloudStack, and can not be modified by the user. This name conforms to the requirements in IETF RFC 1123.
- Display name – the name displayed in the CloudStack web UI. Can be set by the user. Defaults to instance name.
- Name – host name that the DHCP server assigns to the VM. Can be set by the user. Defaults to instance name

Guest VMs can be configured to be Highly Available (HA). An HA-enabled VM is monitored by the system. If the system detects that the VM is down, it will attempt to restart the VM, possibly on a different host. For more information, see HA-Enabled Virtual Machines on

Each new VM is allocated one public IP address. When the VM is started, CloudStack automatically creates a static NAT between this public IP address and the private IP address of the VM.

If elastic IP is in use (with the NetScaler load balancer), the IP address initially allocated to the new VM is not marked as elastic. The user must replace the automatically configured IP with a specifically acquired elastic IP, and set up the static NAT mapping between this new IP and the guest VM's private IP. The VM's original IP address is then released and returned to the pool of available public IPs.


CloudStack cannot distinguish a guest VM that was shut down by the user (such as with the "shutdown" command in Linux) from a VM that shut down unexpectedly. If an HA-enabled VM is shut down from inside the VM, CloudStack will restart it. To shut down an HA-enabled VM, you must go through the CloudStack UI or API.

10.2. Best Practices for Virtual Machines

The CloudStack administrator should monitor the total number of VM instances in each cluster, and disable allocation to the cluster if the total is approaching the maximum that the hypervisor can handle. Be sure to leave a safety margin to allow for the possibility of one or more hosts failing, which would increase the VM load on the other hosts as the VMs are automatically redeployed. Consult the documentation for your chosen hypervisor to find the maximum permitted number of VMs per host, then use CloudStack global configuration settings to set this as the default limit. Monitor the VM activity in each cluster at all times. Keep the total number of VMs below a safe level that allows for the occasional host failure. For example, if there are N hosts in the cluster, and you want to allow for one host in the cluster to be down at any given time, the total number of VM instances you can permit in the cluster is at most $(N-1) * (\text{per-host-limit})$. Once a cluster reaches this number of VMs, use the CloudStack UI to disable allocation of more VMs to the cluster.

10.3. VM Lifecycle

Virtual machines can be in the following states:



Once a virtual machine is destroyed, it cannot be recovered. All the resources used by the virtual machine will be reclaimed by the system. This includes the virtual machine's IP address.

A stop will attempt to gracefully shut down the operating system, which typically involves terminating all the running applications. If the operation system cannot be stopped, it will be forcefully terminated. This has the same effect as pulling the power cord to a physical machine.

A reboot is a stop followed by a start.

CloudStack preserves the state of the virtual machine hard disk until the machine is destroyed.

A running virtual machine may fail because of hardware or network issues. A failed virtual machine is in the down state.

The system places the virtual machine into the down state if it does not receive the heartbeat from the hypervisor for three minutes.

The user can manually restart the virtual machine from the down state.

The system will start the virtual machine from the down state automatically if the virtual machine is marked as HA-enabled.

10.4. Creating VMs

Virtual machines are usually created from a template. Users can also create blank virtual machines. A blank virtual machine is a virtual machine without an OS template. Users can attach an ISO file and install the OS from the CD/DVD-ROM.

**Nota**

You can create a VM without starting it. You can determine whether the VM needs to be started as part of the VM deployment. A request parameter, `startVM`, in the `deployVm` API provides this feature. For more information, see the [Developer's Guide](#)

To create a VM from a template:

1. Log in to the CloudStack UI as an administrator or user.
2. In the left navigation bar, click Instances.
3. Click Add Instance.
4. Select a zone.
5. Select a template, then follow the steps in the wizard. For more information about how the templates came to be in this list, see [Capítulo 12, Trabalhando com templates](#).
6. Be sure that the hardware you have allows starting the selected service offering.
7. Click Submit and your VM will be created and started.

**Nota**

For security reason, the internal name of the VM is visible only to the root admin.

To create a VM from an ISO:

**Nota**

(XenServer) Windows VMs running on XenServer require PV drivers, which may be provided in the template or added after the VM is created. The PV drivers are necessary for essential management functions such as mounting additional volumes and ISO images, live migration, and graceful shutdown.

1. Log in to the CloudStack UI as an administrator or user.
2. In the left navigation bar, click Instances.
3. Click Add Instance.
4. Select a zone.
5. Select ISO Boot, and follow the steps in the wizard.


6. Click Submit and your VM will be created and started.

10.5. Accessing VMs

Any user can access their own virtual machines. The administrator can access all VMs running in the cloud.

To access a VM through the CloudStack UI:

1. Log in to the CloudStack UI as a user or admin.
2. Click Instances, then click the name of a running VM.
3. Click the View Console button



To access a VM directly over the network:

1. The VM must have some port open to incoming traffic. For example, in a basic zone, a new VM might be assigned to a security group which allows incoming traffic. This depends on what security group you picked when creating the VM. In other cases, you can open a port by setting up a port forwarding policy. See [Seção 15.14, “Encaminhamento de IP e firewall”](#).
2. If a port is open but you can not access the VM using ssh, it's possible that ssh is not already enabled on the VM. This will depend on whether ssh is enabled in the template you picked when creating the VM. Access the VM through the CloudStack UI and enable ssh on the machine using the commands for the VM's operating system.
3. If the network has an external firewall device, you will need to create a firewall rule to allow access. See [Seção 15.14, “Encaminhamento de IP e firewall”](#).



10.6. Stopping and Starting VMs

Once a VM instance is created, you can stop, restart, or delete it as needed. In the CloudStack UI, click Instances, select the VM, and use the Stop, Start, Reboot, and Destroy links.

10.7. Changing the VM Name, OS, or Group



After a VM is created, you can modify the display name, operating system, and the group it belongs to.

To access a VM through the CloudStack UI:

1. Log in to the CloudStack UI as a user or admin.
2. In the left navigation, click Instances.
3. Select the VM that you want to modify.
4. Click the Stop button to stop the VM. 
5. Click Edit. 
6. Make the desired changes to the following:
 7. **Display name:** Enter a new display name if you want to change the name of the VM.
 8. **OS Type:** Select the desired operating system.
 9. **Group:** Enter the group name for the VM.
10. Click Apply.

10.8. Changing the Service Offering for a VM

To upgrade or downgrade the level of compute resources available to a virtual machine, you can change the VM's compute offering.

1. Log in to the CloudStack UI as a user or admin.
2. In the left navigation, click Instances.
3. Choose the VM that you want to work with.
4. Click the Stop button to stop the VM. 
5. Click the Change Service button. 

The Change service dialog box is displayed.
6. Select the offering you want to apply to the selected VM.
7. Click OK.


10.9. Moving VMs Between Hosts (Manual Live Migration)

The CloudStack administrator can move a running VM from one host to another without interrupting service to users or going into maintenance mode. This is called manual live migration, and can be done under the following conditions:

- The root administrator is logged in. Domain admins and users can not perform manual live migration of VMs.

- The VM is running. Stopped VMs can not be live migrated.
- The destination host must be in the same cluster as the original host.
- The VM must not be using local disk storage.
- The destination host must have enough available capacity. If not, the VM will remain in the "migrating" state until memory becomes available.


To manually live migrate a virtual machine

1. Log in to the CloudStack UI as a user or admin.
2. In the left navigation, click Instances.
3. Choose the VM that you want to migrate.
4. Click the Migrate Instance button. 
5. From the list of hosts, choose the one to which you want to move the VM.
6. Clique em OK.

10.10. Deleting VMs

Users can delete their own virtual machines. A running virtual machine will be abruptly stopped before it is deleted. Administrators can delete any virtual machines.

To delete a virtual machine:

1. Log in to the CloudStack UI as a user or admin.
2. In the left navigation, click Instances.
3. Choose the VM that you want to delete.
4. Click the Destroy Instance button. 

10.11. Working with ISOs

CloudStack supports ISOs and their attachment to guest VMs. An ISO is a read-only file that has an ISO/CD-ROM style file system. Users can upload their own ISOs and mount them on their guest VMs.

ISOs are uploaded based on a URL. HTTP is the supported protocol. Once the ISO is available via HTTP specify an upload URL such as `http://my.web.server/filename.iso`.

ISOs may be public or private, like templates. ISOs are not hypervisor-specific. That is, a guest on vSphere can mount the exact same image that a guest on KVM can mount.

ISO images may be stored in the system and made available with a privacy level similar to templates. ISO images are classified as either bootable or not bootable. A bootable ISO image is one that contains an OS image. CloudStack allows a user to boot a guest VM off of an ISO image. Users can also attach ISO images to guest VMs. For example, this enables installing PV drivers into Windows. ISO images are not hypervisor-specific.

10.11.1. Adding an ISO

To make additional operating system or other software available for use with guest VMs, you can add an ISO. The ISO is typically thought of as an operating system image, but you can also add ISOs for other types of software, such as desktop applications that you want to be installed as part of a template.

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. In the left navigation bar, click Templates.
3. In Select View, choose ISOs.
4. Click Add ISO.
5. In the Add ISO screen, provide the following:
 - **Name:** Short name for the ISO image. For example, CentOS 6.2 64-bit.
 - **Description:** Display text for the ISO image. For example, CentOS 6.2 64-bit.
 - **URL:** The URL that hosts the ISO image. The Management Server must be able to access this location via HTTP. If needed you can place the ISO image directly on the Management Server
 - **Zone:** Choose the zone where you want the ISO to be available, or All Zones to make it available throughout CloudStack.
 - **Bootable:** Whether or not a guest could boot off this ISO image. For example, a CentOS ISO is bootable, a Microsoft Office ISO is not bootable.
 - **OS Type:** This helps CloudStack and the hypervisor perform certain operations and make assumptions that improve the performance of the guest. Select one of the following.
 - If the operating system of your desired ISO image is listed, choose it.
 - If the OS Type of the ISO is not listed or if the ISO is not bootable, choose Other.
 - (XenServer only) If you want to boot from this ISO in PV mode, choose Other PV (32-bit) or Other PV (64-bit)
 - (KVM only) If you choose an OS that is PV-enabled, the VMs created from this ISO will have a SCSI (virtio) root disk. If the OS is not PV-enabled, the VMs will have an IDE root disk. The PV-enabled types are:

Fedora 13	Fedora 12	Fedora 11
Fedora 10	Fedora 9	Other PV
Debian GNU/Linux	CentOS 5.3	CentOS 5.4
CentOS 5.5	Red Hat Enterprise Linux 5.3	Red Hat Enterprise Linux 5.4
Red Hat Enterprise Linux 5.5	Red Hat Enterprise Linux 6	




Nota

It is not recommended to choose an older version of the OS than the version in the image. For example, choosing CentOS 5.4 to support a CentOS 6.2 image will usually not work. In these cases, choose Other.

- **Extractable:** Choose Yes if the ISO should be available for extraction.
 - **Public:** Choose Yes if this ISO should be available to other users.
 - **Featured:** Choose Yes if you would like this ISO to be more prominent for users to select. The ISO will appear in the Featured ISOs list. Only an administrator can make an ISO Featured.
6. Clique em OK.
- The Management Server will download the ISO. Depending on the size of the ISO, this may take a long time. The ISO status column will display Ready once it has been successfully downloaded into secondary storage. Clicking Refresh updates the download percentage.
7. **Important:** Wait for the ISO to finish downloading. If you move on to the next task and try to use the ISO right away, it will appear to fail. The entire ISO must be available before CloudStack can work with it.

10.11.2. Attaching an ISO to a VM

1. In the left navigation, click Instances.
2. Choose the virtual machine you want to work with.
3. Click the Attach ISO button. 
4. In the Attach ISO dialog box, select the desired ISO.
5. Clique em OK.

Working With Hosts

11.1. Adding Hosts

Additional hosts can be added at any time to provide more capacity for guest VMs. For requirements and instructions, see [Seção 7.6, "Adding a Host"](#).

11.2. Scheduled Maintenance and Maintenance Mode for Hosts

You can place a host into maintenance mode. When maintenance mode is activated, the host becomes unavailable to receive new guest VMs, and the guest VMs already running on the host are seamlessly migrated to another host not in maintenance mode. This migration uses live migration technology and does not interrupt the execution of the guest.

11.2.1. vCenter and Maintenance Mode

To enter maintenance mode on a vCenter host, both vCenter and CloudStack must be used in concert. CloudStack and vCenter have separate maintenance modes that work closely together.

1. Place the host into CloudStack's "scheduled maintenance" mode. This does not invoke the vCenter maintenance mode, but only causes VMs to be migrated off the host

When the CloudStack maintenance mode is requested, the host first moves into the Prepare for Maintenance state. In this state it cannot be the target of new guest VM starts. Then all VMs will be migrated off the server. Live migration will be used to move VMs off the host. This allows the guests to be migrated to other hosts with no disruption to the guests. After this migration is completed, the host will enter the Ready for Maintenance mode.

2. Wait for the "Ready for Maintenance" indicator to appear in the UI.
3. Now use vCenter to perform whatever actions are necessary to maintain the host. During this time, the host cannot be the target of new VM allocations.
4. When the maintenance tasks are complete, take the host out of maintenance mode as follows:

- a. First use vCenter to exit the vCenter maintenance mode.

This makes the host ready for CloudStack to reactivate it.

- b. Then use CloudStack's administrator UI to cancel the CloudStack maintenance mode

When the host comes back online, the VMs that were migrated off of it may be migrated back to it manually and new VMs can be added.

11.2.2. XenServer and Maintenance Mode

For XenServer, you can take a server offline temporarily by using the Maintenance Mode feature in XenCenter. When you place a server into Maintenance Mode, all running VMs are automatically migrated from it to another host in the same pool. If the server is the pool master, a new master will also be selected for the pool. While a server is Maintenance Mode, you cannot create or start any VMs on it.

To place a server in Maintenance Mode:

1. In the Resources pane, select the server, then do one of the following:
 - Right-click, then click Enter Maintenance Mode on the shortcut menu.
 - On the Server menu, click Enter Maintenance Mode.
2. Click Enter Maintenance Mode.

The server's status in the Resources pane shows when all running VMs have been successfully migrated off the server.



To take a server out of Maintenance Mode:

1. In the Resources pane, select the server, then do one of the following:
 - Right-click, then click Exit Maintenance Mode on the shortcut menu.
 - On the Server menu, click Exit Maintenance Mode.
2. Click Exit Maintenance Mode.

11.3. Disabling and Enabling Zones, Pods, and Clusters

You can enable or disable a zone, pod, or cluster without permanently removing it from the cloud. This is useful for maintenance or when there are problems that make a portion of the cloud infrastructure unreliable. No new allocations will be made to a disabled zone, pod, or cluster until its state is returned to Enabled. When a zone, pod, or cluster is first added to the cloud, it is Disabled by default.

To disable and enable a zone, pod, or cluster:

1. Log in to the CloudStack UI as administrator
2. In the left navigation bar, click Infrastructure.
3. In Zones, click View More.
4. If you are disabling or enabling a zone, find the name of the zone in the list, and click the Enable/Disable button. 
5. If you are disabling or enabling a pod or cluster, click the name of the zone that contains the pod or cluster.
6. Click the Compute tab.
7. In the Pods or Clusters node of the diagram, click View All.
8. Click the pod or cluster name in the list.
9.  Click the Enable/Disable button.

11.4. Removing Hosts

Hosts can be removed from the cloud as needed. The procedure to remove a host depends on the hypervisor type.

11.4.1. Removing XenServer and KVM Hosts

A node cannot be removed from a cluster until it has been placed in maintenance mode. This will ensure that all of the VMs on it have been migrated to other Hosts. To remove a Host from the cloud:

1. Place the node in maintenance mode.

See [Seção 11.2, “Scheduled Maintenance and Maintenance Mode for Hosts”](#).

2. For KVM, stop the cloudstack-agent service.
3. Use the UI option to remove the node.

Then you may power down the Host, re-use its IP address, re-install it, etc

11.4.2. Removing vSphere Hosts

To remove this type of host, first place it in maintenance mode, as described in [Seção 11.2, “Scheduled Maintenance and Maintenance Mode for Hosts”](#). Then use CloudStack to remove the host. CloudStack will not direct commands to a host that has been removed using CloudStack. However, the host may still exist in the vCenter cluster.

11.5. Re-Installing Hosts

You can re-install a host after placing it in maintenance mode and then removing it. If a host is down and cannot be placed in maintenance mode, it should still be removed before the re-install.

11.6. Mantendo hipervisores em hosts

Ao executar software de hipervisor em hosts, certifique-se de que todas as hotfixes providas pelo fornecedor do hipervisor estejam aplicadas. Acompanhe a liberação de correções do hipervisor através do canal de suporte do fornecedor, e aplique as correções assim que possível após sua liberação. O CloudStack não acompanhará ou notificará você sobre correções requeridas no hipervisor. É essencial que seus hosts estejam completamente atualizados com as correções fornecidas para o hipervisor. O fornecedor do hipervisor provavelmente recusará suporte a qualquer sistema que não esteja atualizado com as correções fornecidas.



Nota

A falta de hotfixes atualizadas pode levar a dados corrompidos e à perda de máquinas virtuais.

(XenServer) For more information, see [Highly Recommended Hotfixes for XenServer in the CloudStack Knowledge Base](#)¹.

11.7. Changing Host Password

The password for a XenServer Node, KVM Node, or vSphere Node may be changed in the database. Note that all Nodes in a Cluster must have the same password.

¹ http://docs.cloudstack.org/Knowledge_Base/Possible_VM_corruption_if_XenServer_Hotfix_is_not_Applied/Highly_Recommended_Hotfixes_for_XenServer_5.6_SP2

To change a Node's password:

1. Identify all hosts in the cluster.
2. Change the password on all hosts in the cluster. Now the password for the host and the password known to CloudStack will not match. Operations on the cluster will fail until the two passwords match.
3. Get the list of host IDs for the host in the cluster where you are changing the password. You will need to access the database to determine these host IDs. For each hostname "h" (or vSphere cluster) that you are changing the password for, execute:

```
mysql> select id from cloud.host where name like '%h%';
```

4. This should return a single ID. Record the set of such IDs for these hosts.
5. Update the passwords for the host in the database. In this example, we change the passwords for hosts with IDs 5, 10, and 12 to "password".

```
mysql> update cloud.host set password='password' where id=5 or id=10 or id=12;
```

11.8. Host Allocation

The system automatically picks the most appropriate host to run each virtual machine. End users may specify the zone in which the virtual machine will be created. End users do not have control over which host will run the virtual machine instance.

CloudStack administrators can specify that certain hosts should have a preference for particular types of guest instances. For example, an administrator could state that a host should have a preference to run Windows guests. The default host allocator will attempt to place guests of that OS type on such hosts first. If no such host is available, the allocator will place the instance wherever there is sufficient physical capacity.

Both vertical and horizontal allocation is allowed. Vertical allocation consumes all the resources of a given host before allocating any guests on a second host. This reduces power consumption in the cloud. Horizontal allocation places a guest on each host in a round-robin fashion. This may yield better performance to the guests in some cases. CloudStack also allows an element of CPU over-provisioning as configured by the administrator. Over-provisioning allows the administrator to commit more CPU cycles to the allocated guests than are actually available from the hardware.

CloudStack also provides a pluggable interface for adding new allocators. These custom allocators can provide any policy the administrator desires.

11.8.1. Over-Provisioning and Service Offering Limits

CloudStack performs CPU over-provisioning based on an over-provisioning ratio configured by the administrator. This is defined by the `cpu.overprovisioning.factor` global configuration variable.

CloudStack performs CPU over-provisioning based on an over-provisioning ratio configured by the administrator. This is defined by the `cpu.overprovisioning.factor` global configuration variable

Service offerings limits (e.g. 1 GHz, 1 core) are strictly enforced for core count. For example, a guest with a service offering of one core will have only one core available to it regardless of other activity on the Host.

Service offering limits for gigahertz are enforced only in the presence of contention for CPU resources. For example, suppose that a guest was created with a service offering of 1 GHz on a Host that has 2 GHz cores, and that guest is the only guest running on the Host. The guest will have the full 2 GHz available to it. When multiple guests are attempting to use the CPU a weighting factor is used to schedule CPU resources. The weight is based on the clock speed in the service offering. Guests receive a CPU allocation that is proportionate to the GHz in the service offering. For example, a guest created from a 2 GHz service offering will receive twice the CPU allocation as a guest created from a 1 GHz service offering. CloudStack does not perform memory over-provisioning.

11.9. VLAN Provisioning

CloudStack automatically creates and destroys interfaces bridged to VLANs on the hosts. In general the administrator does not need to manage this process.

CloudStack manages VLANs differently based on hypervisor type. For XenServer or KVM, the VLANs are created on only the hosts where they will be used and then they are destroyed when all guests that require them have been terminated or moved to another host.

For vSphere the VLANs are provisioned on all hosts in the cluster even if there is no guest running on a particular Host that requires the VLAN. This allows the administrator to perform live migration and other functions in vCenter without having to create the VLAN on the destination Host. Additionally, the VLANs are not removed from the Hosts when they are no longer needed.

You can use the same VLANs on different physical networks provided that each physical network has its own underlying layer-2 infrastructure, such as switches. For example, you can specify VLAN range 500 to 1000 while deploying physical networks A and B in an Advanced zone setup. This capability allows you to set up an additional layer-2 physical infrastructure on a different physical NIC and use the same set of VLANs if you run out of VLANs. Another advantage is that you can use the same set of IPs for different customers, each one with their own routers and the guest networks on different physical NICs.

Trabalhando com templates

Um template é uma configuração reusável para máquinas virtuais. Quando usuários criam máquinas virtuais, eles podem escolher em uma lista de templates no CloudStack.

Especificamente, um template é uma imagem de disco virtual que inclui um sistema operacional, software adicional opcional, tais como aplicativos de escritório, e customizações tais como controle de acesso para determinar quem pode utilizar o template. Cada template é associado com um tipo específico de hipervisor, que é especificado quando o template é adicionado ao CloudStack.

O CloudStack vem com um template default. Para apresentar mais opções para os usuários, os administradores e usuários do CloudStack podem criar templates e adicioná-los ao CloudStack.

12.1. Creating Templates: Overview

CloudStack ships with a default template for the CentOS operating system. There are a variety of ways to add more templates. Administrators and end users can add templates. The typical sequence of events is:

1. Launch a VM instance that has the operating system you want. Make any other desired configuration changes to the VM.
2. Stop the VM.
3. Convert the volume into a template.

There are other ways to add templates to CloudStack. For example, you can take a snapshot of the VM's volume and create a template from the snapshot, or import a VHD from another system into CloudStack.

The various techniques for creating templates are described in the next few sections.

12.2. Requirements for Templates

- For XenServer, install PV drivers / Xen tools on each template that you create. This will enable live migration and clean guest shutdown.
- For vSphere, install VMware Tools on each template that you create. This will enable console view to work properly.

12.3. Best Practices for Templates

If you plan to use large templates (100 GB or larger), be sure you have a 10-gigabit network to support the large templates. A slower network can lead to timeouts and other errors when large templates are used.

12.4. The Default Template

CloudStack includes a CentOS template. This template is downloaded by the Secondary Storage VM after the primary and secondary storage are configured. You can use this template in your production deployment or you can delete it and use custom templates.

The root password for the default template is "password".

Capítulo 12. Trabalhando com templates

A default template is provided for each of XenServer, KVM, and vSphere. The templates that are downloaded depend on the hypervisor type that is available in your cloud. Each template is approximately 2.5 GB physical size.

The default template includes the standard iptables rules, which will block most access to the template excluding ssh.

```
# iptables --list
Chain INPUT (policy ACCEPT)
target     prot opt source                destination
RH-Firewall-1-INPUT  all  --  anywhere                anywhere

Chain FORWARD (policy ACCEPT)
target     prot opt source                destination
RH-Firewall-1-INPUT  all  --  anywhere                anywhere

Chain OUTPUT (policy ACCEPT)
target     prot opt source                destination

Chain RH-Firewall-1-INPUT (2 references)
target     prot opt source                destination
ACCEPT     all  --  anywhere                anywhere
ACCEPT     icmp --  anywhere                anywhere    icmp any
ACCEPT     esp  --  anywhere                anywhere
ACCEPT     ah   --  anywhere                anywhere
ACCEPT     udp  --  anywhere                224.0.0.251    udp dpt:mdns
ACCEPT     udp  --  anywhere                anywhere        udp dpt:ipp
ACCEPT     tcp  --  anywhere                anywhere        tcp dpt:ipp
ACCEPT     all  --  anywhere                anywhere        state RELATED,ESTABLISHED
ACCEPT     tcp  --  anywhere                anywhere        state NEW tcp dpt:ssh
REJECT     all  --  anywhere                anywhere        reject-with icmp-host-
```

12.5. Private and Public Templates

When a user creates a template, it can be designated private or public.

Private templates are only available to the user who created them. By default, an uploaded template is private.

When a user marks a template as “public,” the template becomes available to all users in all accounts in the user’s domain, as well as users in any other domains that have access to the Zone where the template is stored. This depends on whether the Zone, in turn, was defined as private or public. A private Zone is assigned to a single domain, and a public Zone is accessible to any domain. If a public template is created in a private Zone, it is available only to users in the domain assigned to that Zone. If a public template is created in a public Zone, it is available to all users in all domains.

12.6. Creating a Template from an Existing Virtual Machine

Once you have at least one VM set up in the way you want, you can use it as the prototype for other VMs.

1. Create and start a virtual machine using any of the techniques given in [Seção 10.4, “Creating VMs”](#).
2. Make any desired configuration changes on the running VM, then click Stop.
3. Wait for the VM to stop. When the status shows Stopped, go to the next step.
4. Click Create Template and provide the following:

- **Name and Display Text.** These will be shown in the UI, so choose something descriptive.
- **OS Type.** This helps CloudStack and the hypervisor perform certain operations and make assumptions that improve the performance of the guest. Select one of the following.
 - If the operating system of the stopped VM is listed, choose it.
 - If the OS type of the stopped VM is not listed, choose Other.
 - If you want to boot from this template in PV mode, choose Other PV (32-bit) or Other PV (64-bit). This choice is available only for XenServer:



Nota

Note: Generally you should not choose an older version of the OS than the version in the image. For example, choosing CentOS 5.4 to support a CentOS 6.2 image will in general not work. In those cases you should choose Other.

- **Public.** Choose Yes to make this template accessible to all users of this CloudStack installation. The template will appear in the Community Templates list. See [Seção 12.5, “Private and Public Templates”](#).
- **Password Enabled.** Choose Yes if your template has the CloudStack password change script installed. See [Seção 12.13, “Adding Password Management to Your Templates”](#).

5. Clique em Add.

The new template will be visible in the Templates section when the template creation process has been completed. The template is then available when creating a new VM.

12.7. Creating a Template from a Snapshot

If you do not want to stop the VM in order to use the Create Template menu item (as described in [Seção 12.6, “Creating a Template from an Existing Virtual Machine”](#)), you can create a template directly from any snapshot through the CloudStack UI.

12.8. Uploading Templates



vSphere Templates and ISOs

If you are uploading a template that was created using vSphere Client, be sure the OVA file does not contain an ISO. If it does, the deployment of VMs from the template will fail.

Templates are uploaded based on a URL. HTTP is the supported access protocol. Templates are frequently large files. You can optionally gzip them to decrease upload times.

To upload a template:

1. In the left navigation bar, click Templates.
2. Click Register Template.
3. Provide the following:
 - **Name and Description.** These will be shown in the UI, so choose something descriptive.
 - **URL.** The Management Server will download the file from the specified URL, such as `http://my.web.server/filename.vhd.gz`.
 - **Zone.** Choose the zone where you want the template to be available, or All Zones to make it available throughout CloudStack.
 - **OS Type:** This helps CloudStack and the hypervisor perform certain operations and make assumptions that improve the performance of the guest. Select one of the following:
 - If the operating system of the stopped VM is listed, choose it.
 - If the OS type of the stopped VM is not listed, choose Other.



Nota

You should not choose an older version of the OS than the version in the image. For example, choosing CentOS 5.4 to support a CentOS 6.2 image will in general not work. In those cases you should choose Other.

- **Hypervisor:** The supported hypervisors are listed. Select the desired one.
- **Format.** The format of the template upload file, such as VHD or OVA.
- **Password Enabled.** Choose Yes if your template has the CloudStack password change script installed. See [Adding Password Management to Your Templates](#)
- **Extractable.** Choose Yes if the template is available for extraction. If this option is selected, end users can download a full image of a template.
- **Public.** Choose Yes to make this template accessible to all users of this CloudStack installation. The template will appear in the Community Templates list. See [Seção 12.5, “Private and Public Templates”](#).
- **Featured.** Choose Yes if you would like this template to be more prominent for users to select. The template will appear in the Featured Templates list. Only an administrator can make a template Featured.

12.9. Exporting Templates

End users and Administrators may export templates from the CloudStack. Navigate to the template in the UI and choose the Download function from the Actions menu.

12.10. Creating a Windows Template

Windows templates must be prepared with Sysprep before they can be provisioned on multiple machines. Sysprep allows you to create a generic Windows template and avoid any possible SID conflicts.



Nota

(XenServer) Windows VMs running on XenServer require PV drivers, which may be provided in the template or added after the VM is created. The PV drivers are necessary for essential management functions such as mounting additional volumes and ISO images, live migration, and graceful shutdown.

An overview of the procedure is as follows:

1. Upload your Windows ISO.

For more information, see [Seção 10.11.1, “Adding an ISO”](#).

2. Create a VM Instance with this ISO.

For more information, see [Seção 10.4, “Creating VMs”](#).

3. Follow the steps in Sysprep for Windows Server 2008 R2 (below) or Sysprep for Windows Server 2003 R2, depending on your version of Windows Server
4. The preparation steps are complete. Now you can actually create the template as described in [Creating the Windows Template](#).

12.10.1. System Preparation for Windows Server 2008 R2

For Windows 2008 R2, you run Windows System Image Manager to create a custom sysprep response XML file. Windows System Image Manager is installed as part of the Windows Automated Installation Kit (AIK). Windows AIK can be downloaded from [Microsoft Download Center](#)¹.

Use the following steps to run sysprep for Windows 2008 R2:



Nota

The steps outlined here are derived from the excellent guide by Charity Shelbourne, originally published at [Windows Server 2008 Sysprep Mini-Setup](#).²

1. Download and install the Windows AIK

¹ <http://www.microsoft.com/en-us/download/details.aspx?id=9085>

² <http://blogs.technet.com/askcore/archive/2008/10/31/automating-the-oobe-process-during-windows-server-2008-sysprep-mini-setup.aspx>




Nota

Windows AIK should not be installed on the Windows 2008 R2 VM you just created. Windows AIK should not be part of the template you create. It is only used to create the sysprep answer file.


2. Copy the install.wim file in the \sources directory of the Windows 2008 R2 installation DVD to the hard disk. This is a very large file and may take a long time to copy. Windows AIK requires the WIM file to be writable.
3. Start the Windows System Image Manager, which is part of the Windows AIK.
4. In the Windows Image pane, right click the Select a Windows image or catalog file option to load the install.wim file you just copied.
5. Select the Windows 2008 R2 Edition.

You may be prompted with a warning that the catalog file cannot be opened. Click Yes to create a new catalog file.


6. In the Answer File pane, right click to create a new answer file.
7. Generate the answer file from the Windows System Image Manager using the following steps:
 - a. The first page you need to automate is the Language and Country or Region Selection page. To automate this, expand Components in your Windows Image pane, right-click and add the Microsoft-Windows-International-Core setting to Pass 7 oobeSystem. In your Answer File pane, configure the InputLocale, SystemLocale, UILanguage, and UserLocale with the appropriate settings for your language and country or region. Should you have a question about any of these settings, you can right-click on the specific setting and select Help. This will open the appropriate CHM help file with more information, including examples on the setting you are attempting to configure.



- b. You need to automate the Software License Terms Selection page, otherwise known as the End-User License Agreement (EULA). To do this, expand the Microsoft-Windows-Shell-Setup component. High-light the OOBE setting, and add the setting to the Pass 7 oobeSystem. In Settings, set HideEULAPage true.



- c. Make sure the license key is properly set. If you use MAK key, you can just enter the MAK key on the Windows 2008 R2 VM. You need not input the MAK into the Windows System Image Manager. If you use KMS host for activation you need not enter the Product Key. Details of Windows Volume Activation can be found at <http://technet.microsoft.com/en-us/library/bb892849.aspx>
- d. You need to automate is the Change Administrator Password page. Expand the Microsoft-Windows-Shell-Setup component (if it is not still expanded), expand UserAccounts, right-click on AdministratorPassword, and add the setting to the Pass 7 oobeSystem configuration pass of your answer file. Under Settings, specify a password next to Value.



You may read the AIK documentation and set many more options that suit your deployment. The steps above are the minimum needed to make Windows unattended setup work.

8. Save the answer file as unattend.xml. You can ignore the warning messages that appear in the validation window.
9. Copy the unattend.xml file into the c:\windows\system32\sysprep directory of the Windows 2008 R2 Virtual Machine
10. Once you place the unattend.xml file in c:\windows\system32\sysprep directory, you run the sysprep tool as follows:

```
cd c:\Windows\System32\sysprep
sysprep.exe /oobe /generalize /shutdown
```

The Windows 2008 R2 VM will automatically shut down after sysprep is complete.

12.10.2. System Preparation for Windows Server 2003 R2

Earlier versions of Windows have a different sysprep tool. Follow these steps for Windows Server 2003 R2.

1. Extract the content of \support\tools\deploy.cab on the Windows installation CD into a directory called c:\sysprep on the Windows 2003 R2 VM.
2. Run c:\sysprep\setupmgr.exe to create the sysprep.inf file.

- a. Select Create New to create a new Answer File.
 - b. Enter “Sysprep setup” for the Type of Setup.
 - c. Select the appropriate OS version and edition.
 - d. On the License Agreement screen, select “Yes fully automate the installation”.
 - e. Provide your name and organization.
 - f. Leave display settings at default.
 - g. Set the appropriate time zone.
 - h. Provide your product key.
 - i. Select an appropriate license mode for your deployment
 - j. Select “Automatically generate computer name”.
 - k. Type a default administrator password. If you enable the password reset feature, the users will not actually use this password. This password will be reset by the instance manager after the guest boots up.
 - l. Leave Network Components at “Typical Settings”.
 - m. Select the “WORKGROUP” option.
 - n. Leave Telephony options at default.
 - o. Select appropriate Regional Settings.
 - p. Select appropriate language settings.
 - q. Do not install printers.
 - r. Do not specify “Run Once commands”.
 - s. You need not specify an identification string.
 - t. Save the Answer File as c:\sysprep\sysprep.inf.
3. Run the following command to sysprep the image:

```
c:\sysprep\sysprep.exe -reseal -mini -activated
```

After this step the machine will automatically shut down

12.11. Importing Amazon Machine Images

The following procedures describe how to import an Amazon Machine Image (AMI) into CloudStack when using the XenServer hypervisor.

Assume you have an AMI file and this file is called CentOS_6.2_x64. Assume further that you are working on a CentOS host. If the AMI is a Fedora image, you need to be working on a Fedora host initially.

You need to have a XenServer host with a file-based storage repository (either a local ext3 SR or an NFS SR) to convert to a VHD once the image file has been customized on the Centos/Fedora host.



Nota

Ao copiar e colar um comando, certifique-se que o comando tenha colado como uma única linha antes de executar. Alguns viewers documento pode introduzir quebras de linha indesejadas no texto copiado.

To import an AMI:

1. Set up loopback on image file:

```
# mkdir -p /mnt/loop/centos62
# mount -o loop CentOS_6.2_x64 /mnt/loop/centos54
```

2. Install the kernel-xen package into the image. This downloads the PV kernel and ramdisk to the image.

```
# yum -c /mnt/loop/centos54/etc/yum.conf --installroot=/mnt/loop/centos62/ -y install
kernel-xen
```

3. Create a grub entry in /boot/grub/grub.conf.

```
# mkdir -p /mnt/loop/centos62/boot/grub
# touch /mnt/loop/centos62/boot/grub/grub.conf
# echo "" > /mnt/loop/centos62/boot/grub/grub.conf
```

4. Determine the name of the PV kernel that has been installed into the image.

```
# cd /mnt/loop/centos62
# ls lib/modules/
2.6.16.33-xenU 2.6.16-xenU 2.6.18-164.15.1.el5xen 2.6.18-164.6.1.el5.centos.plus
2.6.18-xenU-ec2-v1.0 2.6.21.7-2.fc8xen 2.6.31-302-ec2
# ls boot/initrd*
boot/initrd-2.6.18-164.6.1.el5.centos.plus.img boot/initrd-2.6.18-164.15.1.el5xen.img
# ls boot/vmlinuz*
boot/vmlinuz-2.6.18-164.15.1.el5xen boot/vmlinuz-2.6.18-164.6.1.el5.centos.plus boot/
vmlinuz-2.6.18-xenU-ec2-v1.0 boot/vmlinuz-2.6.21-2952.fc8xen
```

Xen kernels/ramdisk always end with "xen". For the kernel version you choose, there has to be an entry for that version under lib/modules, there has to be an initrd and vmlinuz corresponding to that. Above, the only kernel that satisfies this condition is 2.6.18-164.15.1.el5xen.

5. Based on your findings, create an entry in the grub.conf file. Below is an example entry.

```
default=0
timeout=5
hiddenmenu
title CentOS (2.6.18-164.15.1.el5xen)
    root (hd0,0)
    kernel /boot/vmlinuz-2.6.18-164.15.1.el5xen ro root=/dev/xvda
```

```
initrd /boot/initrd-2.6.18-164.15.1.el5xen.img
```

6. Edit `etc/fstab`, changing “`sda1`” to “`xvda`” and changing “`sdb`” to “`xvdb`”.

```
# cat etc/fstab
/dev/xvda / ext3 defaults 1 1
/dev/xvdb /mnt ext3 defaults 0 0
none /dev/pts devpts gid=5,mode=620 0 0
none /proc proc defaults 0 0
none /sys sysfs defaults 0 0
```

7. Enable login via the console. The default console device in a XenServer system is `xvc0`. Ensure that `etc/inittab` and `etc/securetty` have the following lines respectively:

```
# grep xvc0 etc/inittab
co:2345:respawn:/sbin/agetty xvc0 9600 vt100-nav
# grep xvc0 etc/securetty
xvc0
```

8. Ensure the ramdisk supports PV disk and PV network. Customize this for the kernel version you have determined above.

```
# chroot /mnt/loop/centos54
# cd /boot/
# mv initrd-2.6.18-164.15.1.el5xen.img initrd-2.6.18-164.15.1.el5xen.img.bak
# mkinitrd -f /boot/initrd-2.6.18-164.15.1.el5xen.img --with=xennet --preload=xenblk --omit-scsi-modules 2.6.18-164.15.1.el5xen
```

9. Change the password.

```
# passwd
Changing password for user root.
New UNIX password:
Retype new UNIX password:
passwd: all authentication tokens updated successfully.
```

10. Exit out of `chroot`.

```
# exit
```

11. Check `etc/ssh/sshd_config` for lines allowing ssh login using a password.

```
# egrep "PermitRootLogin|PasswordAuthentication" /mnt/loop/centos54/etc/ssh/sshd_config
PermitRootLogin yes
PasswordAuthentication yes
```

12. If you need the template to be enabled to reset passwords from the CloudStack UI or API, install the password change script into the image at this point. See [Seção 12.13, “Adding Password Management to Your Templates”](#).

13. Unmount and delete loopback mount.

```
# umount /mnt/loop/centos54
# losetup -d /dev/loop0
```


- Copy the image file to your XenServer host's file-based storage repository. In the example below, the Xenserver is "xenhost". This XenServer has an NFS repository whose uuid is a9c5b8c8-536b-a193-a6dc-51af3e5ff799.

```
# scp CentOS_6.2_x64 xenhost:/var/run/sr-mount/a9c5b8c8-536b-a193-a6dc-51af3e5ff799/
```

- Log in to the Xenserver and create a VDI the same size as the image.

```
[root@xenhost ~]# cd /var/run/sr-mount/a9c5b8c8-536b-a193-a6dc-51af3e5ff799
[root@xenhost a9c5b8c8-536b-a193-a6dc-51af3e5ff799]# ls -lh CentOS_6.2_x64
-rw-r--r-- 1 root root 10G Mar 16 16:49 CentOS_6.2_x64
[root@xenhost a9c5b8c8-536b-a193-a6dc-51af3e5ff799]# xe vdi-create virtual-size=10GiB sr-
uuid=a9c5b8c8-536b-a193-a6dc-51af3e5ff799 type=user name-label="Centos 6.2 x86_64"
cad7317c-258b-4ef7-b207-cdf0283a7923
```

- Import the image file into the VDI. This may take 10–20 minutes.

```
[root@xenhost a9c5b8c8-536b-a193-a6dc-51af3e5ff799]# xe vdi-import
filename=CentOS_6.2_x64 uuid=cad7317c-258b-4ef7-b207-cdf0283a7923
```

- Locate a the VHD file. This is the file with the VDI's UUID as its name. Compress it and upload it to your web server.

```
[root@xenhost a9c5b8c8-536b-a193-a6dc-51af3e5ff799]# bzip2 -c cad7317c-258b-4ef7-b207-
cdf0283a7923.vhd > CentOS_6.2_x64.vhd.bz2
[root@xenhost a9c5b8c8-536b-a193-a6dc-51af3e5ff799]# scp CentOS_6.2_x64.vhd.bz2
webserver:/var/www/html/templates/
```

12.12. Converting a Hyper-V VM to a Template

To convert a Hyper-V VM to a XenServer-compatible CloudStack template, you will need a standalone XenServer host with an attached NFS VHD SR. Use whatever XenServer version you are using with CloudStack, but use XenCenter 5.6 FP1 or SP2 (it is backwards compatible to 5.6). Additionally, it may help to have an attached NFS ISO SR.

For Linux VMs, you may need to do some preparation in Hyper-V before trying to get the VM to work in XenServer. Clone the VM and work on the clone if you still want to use the VM in Hyper-V. Uninstall Hyper-V Integration Components and check for any references to device names in `/etc/fstab`:

- From the `linux_ic/drivers/dist` directory, run `make uninstall` (where "linux_ic" is the path to the copied Hyper-V Integration Components files).
- Restore the original `initrd` from backup in `/boot/` (the backup is named `*.backup0`).
- Remove the "hdX=noprobe" entries from `/boot/grub/menu.lst`.
- Check `/etc/fstab` for any partitions mounted by device name. Change those entries (if any) to mount by LABEL or UUID. You can get that information with the `blkid` command.

The next step is make sure the VM is not running in Hyper-V, then get the VHD into XenServer. There are two options for doing this.

Option one:

- Import the VHD using XenCenter. In XenCenter, go to Tools>Virtual Appliance Tools>Disk Image Import.

2. Choose the VHD, then click Next.
3. Name the VM, choose the NFS VHD SR under Storage, enable "Run Operating System Fixups" and choose the NFS ISO SR.
4. Click Next, then Finish. A VM should be created.

Option two:

1. Run XenConvert, under From choose VHD, under To choose XenServer. Click Next.
2. Choose the VHD, then click Next.
3. Input the XenServer host info, then click Next.
4. Name the VM, then click Next, then Convert. A VM should be created.

Once you have a VM created from the Hyper-V VHD, prepare it using the following steps:

1. Boot the VM, uninstall Hyper-V Integration Services, and reboot.
2. Install XenServer Tools, then reboot.
3. Prepare the VM as desired. For example, run sysprep on Windows VMs. See [Seção 12.10, "Creating a Windows Template"](#).

Either option above will create a VM in HVM mode. This is fine for Windows VMs, but Linux VMs may not perform optimally. Converting a Linux VM to PV mode will require additional steps and will vary by distribution.

1. Shut down the VM and copy the VHD from the NFS storage to a web server; for example, mount the NFS share on the web server and copy it, or from the XenServer host use sftp or scp to upload it to the web server.
2. In CloudStack, create a new template using the following values:
 - URL. Give the URL for the VHD
 - OS Type. Use the appropriate OS. For PV mode on CentOS, choose Other PV (32-bit) or Other PV (64-bit). This choice is available only for XenServer.
 - Hypervisor. XenServer
 - Format. VHD

The template will be created, and you can create instances from it.

12.13. Adding Password Management to Your Templates

CloudStack provides an optional password reset feature that allows users to set a temporary admin or root password as well as reset the existing admin or root password from the CloudStack UI.

To enable the Reset Password feature, you will need to download an additional script to patch your template. When you later upload the template into CloudStack, you can specify whether reset admin/root password feature should be enabled for this template.

The password management feature works always resets the account password on instance boot. The script does an HTTP call to the virtual router to retrieve the account password that should be set. As long as the virtual router is accessible the guest will have access to the account password that should

be used. When the user requests a password reset the management server generates and sends a new password to the virtual router for the account. Thus an instance reboot is necessary to effect any password changes.

If the script is unable to contact the virtual router during instance boot it will not set the password but boot will continue normally.

12.13.1. Instalação do sistema operacional Linux

Siga os seguintes passos para iniciar a instalação do Linux:

1. Faça download do script cloud-set-guest-password:

- Linux: <http://cloudstack.org/dl/cloud-set-guest-password>
- Windows: <http://sourceforge.net/projects/cloudstack/files/Password%20Management%20Scripts/CloudInstanceManager.msi/download>

2. Copie este arquivo para /etc/init.d.

Em algumas distribuições Linux, copie o arquivo para /etc/rc.d/init.d.

3. Execute o seguinte comando para tornar executável o script:

```
chmod +x /etc/init.d/cloud-set-guest-password
```

4. Dependendo da distribuição Linux, continue com o passo apropriado.

No Fedora, CentOS/RHEL, e Debian, execute:

```
chkconfig --add cloud-set-guest-password
```

12.13.2. Instalação no Windows

Faça download do instalador, CloudInstanceManager.msi, da [página de download](#)³ e execute o instalador na nova máquina virtual Windows.

12.14. Deleting Templates

Templates may be deleted. In general, when a template spans multiple Zones, only the copy that is selected for deletion will be deleted; the same template in other Zones will not be deleted. The provided CentOS template is an exception to this. If the provided CentOS template is deleted, it will be deleted from all Zones.

When templates are deleted, the VMs instantiated from them will continue to run. However, new VMs cannot be created based on the deleted template.

³ <http://cloudstack.org/download.html>

Working With Storage

13.1. Storage Overview

CloudStack defines two types of storage: primary and secondary. Primary storage can be accessed by either iSCSI or NFS. Additionally, direct attached storage may be used for primary storage. Secondary storage is always accessed using NFS.

There is no ephemeral storage in CloudStack. All volumes on all nodes are persistent.

13.2. Storage primária

Esta seção fornece conceitos e detalhes técnicos sobre a storage primária do CloudStack. Para informações sobre como instalar e configurar storage primária através da interface de usuário do CloudStack, veja o Guia de instalação.

[Seção 2.6, “Sobre storage primária”](#)

13.2.1. Best Practices for Primary Storage

- The speed of primary storage will impact guest performance. If possible, choose smaller, higher RPM drives for primary storage.
- Ensure that nothing is stored on the server. Adding the server to CloudStack will destroy any existing data

13.2.2. Runtime Behavior of Primary Storage

Root volumes are created automatically when a virtual machine is created. Root volumes are deleted when the VM is destroyed. Data volumes can be created and dynamically attached to VMs. Data volumes are not deleted when VMs are destroyed.

Administrators should monitor the capacity of primary storage devices and add additional primary storage as needed. See the Advanced Installation Guide.

Administrators add primary storage to the system by creating a CloudStack storage pool. Each storage pool is associated with a cluster.

13.2.3. Hypervisor Support for Primary Storage

The following table shows storage options and parameters for different hypervisors.

	VMware vSphere	Citrix XenServer	KVM	
Format for Disks, Templates, and Snapshots	VMDK	VHD	QCOW2	
iSCSI support	VMFS	Clustered LVM	Yes, via Shared Mountpoint	
Fiber Channel support	VMFS	Yes, via Existing SR	Yes, via Shared Mountpoint	

	VMware vSphere	Citrix XenServer	KVM	
NFS support	Y	Y	Y	
Local storage support	Y	Y	Y	
Storage over-provisioning	NFS and iSCSI	NFS	NFS	

XenServer uses a clustered LVM system to store VM images on iSCSI and Fiber Channel volumes and does not support over-provisioning in the hypervisor. The storage server itself, however, can support thin-provisioning. As a result the CloudStack can still support storage over-provisioning by running on thin-provisioned storage volumes.

KVM supports "Shared Mountpoint" storage. A shared mountpoint is a file system path local to each server in a given cluster. The path must be the same across all Hosts in the cluster, for example /mnt/primary1. This shared mountpoint is assumed to be a clustered filesystem such as OCFS2. In this case the CloudStack does not attempt to mount or unmount the storage as is done with NFS. The CloudStack requires that the administrator insure that the storage is available

With NFS storage, CloudStack manages the overprovisioning. In this case the global configuration parameter `storage.overprovisioning.factor` controls the degree of overprovisioning. This is independent of hypervisor type.

Local storage is an option for primary storage for vSphere, XenServer, and KVM. When the local disk option is enabled, a local disk storage pool is automatically created on each host. To use local storage for the System Virtual Machines (such as the Virtual Router), set `system.vm.use.local.storage` to true in global configuration.

CloudStack supports multiple primary storage pools in a Cluster. For example, you could provision 2 NFS servers in primary storage. Or you could provision 1 iSCSI LUN initially and then add a second iSCSI LUN when the first approaches capacity.

13.2.4. Storage Tags

Storage may be "tagged". A tag is a text string attribute associated with primary storage, a Disk Offering, or a Service Offering. Tags allow administrators to provide additional information about the storage. For example, that is a "SSD" or it is "slow". Tags are not interpreted by CloudStack. They are matched against tags placed on service and disk offerings. CloudStack requires all tags on service and disk offerings to exist on the primary storage before it allocates root or data disks on the primary storage. Service and disk offering tags are used to identify the requirements of the storage that those offerings have. For example, the high end service offering may require "fast" for its root disk volume.

The interaction between tags, allocation, and volume copying across clusters and pods can be complex. To simplify the situation, use the same set of tags on the primary storage for all clusters in a pod. Even if different devices are used to present those tags, the set of exposed tags can be the same.

13.2.5. Maintenance Mode for Primary Storage

Primary storage may be placed into maintenance mode. This is useful, for example, to replace faulty RAM in a storage device. Maintenance mode for a storage device will first stop any new guests from being provisioned on the storage device. Then it will stop all guests that have any volume on that storage device. When all such guests are stopped the storage device is in maintenance mode and may be shut down. When the storage device is online again you may cancel maintenance mode for

the device. The CloudStack will bring the device back online and attempt to start all guests that were running at the time of the entry into maintenance mode.

13.3. Storage secundária

Esta seção fornece conceitos e detalhes técnicos sobre a storage secundária do CloudStack. Para informações sobre como instalar e configurar storage secundária através da interface de usuário do CloudStack UI, veja o Advanced Installation Guide.

[Seção 2.7, “Sobre storage secundária”](#)

13.4. Working With Volumes

A volume provides storage to a guest VM. The volume can provide for a root disk or an additional data disk. CloudStack supports additional volumes for guest VMs.

Volumes are created for a specific hypervisor type. A volume that has been attached to guest using one hypervisor type (e.g, XenServer) may not be attached to a guest that is using another hypervisor type, for example: vSphere, KVM. This is because the different hypervisors use different disk image formats.

CloudStack defines a volume as a unit of storage available to a guest VM. Volumes are either root disks or data disks. The root disk has "/" in the file system and is usually the boot device. Data disks provide for additional storage, for example: "/opt" or "D:". Every guest VM has a root disk, and VMs can also optionally have a data disk. End users can mount multiple data disks to guest VMs. Users choose data disks from the disk offerings created by administrators. The user can create a template from a volume as well; this is the standard procedure for private template creation. Volumes are hypervisor-specific: a volume from one hypervisor type may not be used on a guest of another hypervisor type.



Nota

CloudStack supports attaching up to 13 data disks to a VM on XenServer hypervisor versions 6.0 and above. For the VMs on other hypervisor types, the data disk limit is 6.

13.4.1. Creating a New Volume

You can add more data disk volumes to a guest VM at any time, up to the limits of your storage capacity. Both CloudStack administrators and users can add volumes to VM instances. When you create a new volume, it is stored as an entity in CloudStack, but the actual storage resources are not allocated on the physical storage device until you attach the volume. This optimization allows the CloudStack to provision the volume nearest to the guest that will use it when the first attachment is made.

13.4.1.1. Using Local Storage for Data Volumes

You can create data volumes on local storage (supported with XenServer, KVM, and VMware). The data volume is placed on the same host as the VM instance that is attached to the data volume. These local data volumes can be attached to virtual machines, detached, re-attached, and deleted just as with the other types of data volume.

Local storage is ideal for scenarios where persistence of data volumes and HA is not required. Some of the benefits include reduced disk I/O latency and cost reduction from using inexpensive local disks.

In order for local volumes to be used, the feature must be enabled for the zone.

You can create a data disk offering for local storage. When a user creates a new VM, they can select this disk offering in order to cause the data disk volume to be placed in local storage.

You can not migrate a VM that has a volume in local storage to a different host, nor migrate the volume itself away to a different host. If you want to put a host into maintenance mode, you must first stop any VMs with local data volumes on that host.

13.4.1.2. To Create a New Volume

1. Log in to the CloudStack UI as a user or admin.
2. In the left navigation bar, click Storage.
3. In Select View, choose Volumes.
4. To create a new volume, click Add Volume, provide the following details, and click OK.
 - Name. Give the volume a unique name so you can find it later.
 - Availability Zone. Where do you want the storage to reside? This should be close to the VM that will use the volume.
 - Disk Offering. Choose the characteristics of the storage.

The new volume appears in the list of volumes with the state “Allocated.” The volume data is stored in CloudStack, but the volume is not yet ready for use

5. To start using the volume, continue to Attaching a Volume

13.4.2. Uploading an Existing Volume to a Virtual Machine

Existing data can be made accessible to a virtual machine. This is called uploading a volume to the VM. For example, this is useful to upload data from a local file system and attach it to a VM. Root administrators, domain administrators, and end users can all upload existing volumes to VMs.

The upload is performed using HTTP. The uploaded volume is placed in the zone's secondary storage

You cannot upload a volume if the preconfigured volume limit has already been reached. The default limit for the cloud is set in the global configuration parameter `max.account.volumes`, but administrators can also set per-domain limits that are different from the global default. See [Setting Usage Limits](#)

To upload a volume:

1. (Optional) Create an MD5 hash (checksum) of the disk image file that you are going to upload. After uploading the data disk, CloudStack will use this value to verify that no data corruption has occurred.
2. Log in to the CloudStack UI as an administrator or user
3. In the left navigation bar, click Storage.
4. Click Upload Volume.
5. Provide the following:
 - Name and Description. Any desired name and a brief description that can be shown in the UI.

- Availability Zone. Choose the zone where you want to store the volume. VMs running on hosts in this zone can attach the volume.
- Format. Choose one of the following to indicate the disk image format of the volume.

Hypervisor	Disk Image Format
XenServer	VHD
VMware	OVA
KVM	QCOW2


- URL. The secure HTTP or HTTPS URL that CloudStack can use to access your disk. The type of file at the URL must match the value chosen in Format. For example, if Format is VHD, the URL might look like the following:

`http://yourFileServerIP/userdata/myDataDisk.vhd`

- MD5 checksum. (Optional) Use the hash that you created in step 1.
6. Wait until the status of the volume shows that the upload is complete. Click Instances - Volumes, find the name you specified in step ???, and make sure the status is Uploaded.

13.4.3. Attaching a Volume

You can attach a volume to a guest VM to provide extra disk storage. Attach a volume when you first create a new volume, when you are moving an existing volume from one VM to another, or after you have migrated a volume from one storage pool to another.

1. Log in to the CloudStack UI as a user or admin.
2. In the left navigation, click Storage.
3. In Select View, choose Volumes.
4. Click the volume name in the Volumes list, then click the Attach Disk button .
5. In the Instance popup, choose the VM to which you want to attach the volume. You will only see instances to which you are allowed to attach volumes; for example, a user will see only instances created by that user, but the administrator will have more choices.
6. When the volume has been attached, you should be able to see it by clicking Instances, the instance name, and View Volumes.

13.4.4. Detaching and Moving Volumes




Nota

This procedure is different from moving disk volumes from one storage pool to another. See VM Storage Migration

A volume can be detached from a guest VM and attached to another guest. Both CloudStack administrators and users can detach volumes from VMs and move them to other VMs.

If the two VMs are in different clusters, and the volume is large, it may take several minutes for the volume to be moved to the new VM.

1. Log in to the CloudStack UI as a user or admin.
2. In the left navigation bar, click Storage, and choose Volumes in Select View. Alternatively, if you know which VM the volume is attached to, you can click Instances, click the VM name, and click View Volumes.
3. Click the name of the volume you want to detach, then click the Detach Disk button. 
4. To move the volume to another VM, follow the steps in [Seção 13.4.3, “Attaching a Volume”](#).

13.4.5. VM Storage Migration

Supported in XenServer, KVM, and VMware.



Nota

This procedure is different from moving disk volumes from one VM to another. See Detaching and Moving Volumes [Seção 13.4.4, “Detaching and Moving Volumes”](#).

You can migrate a virtual machine’s root disk volume or any additional data disk volume from one storage pool to another in the same zone.

You can use the storage migration feature to achieve some commonly desired administration goals, such as balancing the load on storage pools and increasing the reliability of virtual machines by moving them away from any storage pool that is experiencing issues.

13.4.5.1. Migrating a Data Disk Volume to a New Storage Pool

1. Log in to the CloudStack UI as a user or admin.
2. Detach the data disk from the VM. See Detaching and Moving Volumes [Seção 13.4.4, “Detaching and Moving Volumes”](#) (but skip the “reattach” step at the end. You will do that after migrating to new storage).
3. Call the CloudStack API command `migrateVolume` and pass in the volume ID and the ID of any storage pool in the zone.
4. Watch for the volume status to change to Migrating, then back to Ready.
5. Attach the volume to any desired VM running in the same cluster as the new storage server. See Attaching a Volume [Seção 13.4.3, “Attaching a Volume”](#)

13.4.5.2. Migrating a VM Root Volume to a New Storage Pool

When migrating the root disk volume, the VM must first be stopped, and users can not access the VM. After migration is complete, the VM can be restarted.

1. Log in to the CloudStack UI as a user or admin.
2. Detach the data disk from the VM. See Detaching and Moving Volumes [Seção 13.4.4, “Detaching and Moving Volumes”](#) (but skip the “reattach” step at the end. You will do that after migrating to new storage).
3. Stop the VM.
4. Use the CloudStack API command, `migrateVirtualMachine`, with the ID of the VM to migrate and the IDs of a destination host and destination storage pool in the same zone.
5. Watch for the VM status to change to Migrating, then back to Stopped.
6. Restart the VM.

13.4.6. Resizing Volumes

CloudStack provides the ability to resize data disks; CloudStack controls volume size by using disk offerings. This provides CloudStack administrators with the flexibility to choose how much space they want to make available to the end users. Volumes within the disk offerings with the same storage tag can be resized. For example, if you only want to offer 10, 50, and 100 GB offerings, the allowed resize should stay within those limits. That implies if you define a 10 GB, a 50 GB and a 100 GB disk offerings, a user can upgrade from 10 GB to 50 GB, or 50 GB to 100 GB. If you create a custom-sized disk offering, then you have the option to resize the volume by specifying a new, larger size.


Additionally, using the `resizeVolume` API, a data volume can be moved from a static disk offering to a custom disk offering with the size specified. This functionality allows those who might be billing by certain volume sizes or disk offerings to stick to that model, while providing the flexibility to migrate to whatever custom size necessary.


This feature is supported on KVM, XenServer, and VMware hosts. However, shrinking volumes is not supported on VMware hosts.

Before you try to resize a volume, consider the following:

- The VMs associated with the volume are stopped.
- The data disks associated with the volume are removed.
- When a volume is shrunk, the disk associated with it is simply truncated, and doing so would put its content at risk of data loss. Therefore, resize any partitions or file systems before you shrink a data disk so that all the data is moved off from that disk.

To resize a volume:

1. Log in to the CloudStack UI as a user or admin.
2. In the left navigation bar, click Storage.
3. In Select View, choose Volumes.
4. Select the volume name in the Volumes list, then click the Resize Volume button 
5. In the Resize Volume pop-up, choose desired characteristics for the storage.



- a. If you select Custom Disk, specify a custom size.
- b. Click Shrink OK to confirm that you are reducing the size of a volume.

This parameter protects against inadvertent shrinking of a disk, which might lead to the risk of data loss. You must sign off that you know what you are doing.

6. Clique em OK.

13.4.7. Volume Deletion and Garbage Collection

The deletion of a volume does not delete the snapshots that have been created from the volume

When a VM is destroyed, data disk volumes that are attached to the VM are not deleted.

Volumes are permanently destroyed using a garbage collection process. The global configuration variables `expunge.delay` and `expunge.interval` determine when the physical deletion of volumes will occur.

- `expunge.delay`: determines how old the volume must be before it is destroyed, in seconds
- `expunge.interval`: determines how often to run the garbage collection check

Administrators should adjust these values depending on site policies around data retention.

13.5. Working with Snapshots

(Supported for the following hypervisors: **XenServer**, **VMware vSphere**, and **KVM**)

CloudStack supports snapshots of disk volumes. Snapshots are a point-in-time capture of virtual machine disks. Memory and CPU states are not captured.

Snapshots may be taken for volumes, including both root and data disks. The administrator places a limit on the number of stored snapshots per user. Users can create new volumes from the snapshot for recovery of particular files and they can create templates from snapshots to boot from a restored disk.

Users can create snapshots manually or by setting up automatic recurring snapshot policies. Users can also create disk volumes from snapshots, which may be attached to a VM like any other disk volume. Snapshots of both root disks and data disks are supported. However, CloudStack does not currently support booting a VM from a recovered root disk. A disk recovered from snapshot of a root disk is treated as a regular data disk; the data on recovered disk can be accessed by attaching the disk to a VM.

A completed snapshot is copied from primary storage to secondary storage, where it is stored until deleted or purged by newer snapshot.

13.5.1. Snapshot Job Throttling

When a snapshot of a virtual machine is requested, the snapshot job runs on the same host where the VM is running or, in the case of a stopped VM, the host where it ran last. If many snapshots are requested for VMs on a single host, this can lead to problems with too many snapshot jobs overwhelming the resources of the host.

To address this situation, the cloud's root administrator can throttle how many snapshot jobs are executed simultaneously on the hosts in the cloud by using the global configuration setting `concurrent.snapshots.threshold.perhost`. By using this setting, the administrator can better ensure that snapshot jobs do not time out and hypervisor hosts do not experience performance issues due to hosts being overloaded with too many snapshot requests.

Set `concurrent.snapshots.threshold.perhost` to a value that represents a best guess about how many snapshot jobs the hypervisor hosts can execute at one time, given the current resources of the hosts and the number of VMs running on the hosts. If a given host has more snapshot requests, the additional requests are placed in a waiting queue. No new snapshot jobs will start until the number of currently executing snapshot jobs falls below the configured limit.

The admin can also set `job.expire.minutes` to place a maximum on how long a snapshot request will wait in the queue. If this limit is reached, the snapshot request fails and returns an error message.

13.5.2. Automatic Snapshot Creation and Retention

(Supported for the following hypervisors: **XenServer**, **VMware vSphere**, and **KVM**)

Users can set up a recurring snapshot policy to automatically create multiple snapshots of a disk at regular intervals. Snapshots can be created on an hourly, daily, weekly, or monthly interval. One snapshot policy can be set up per disk volume. For example, a user can set up a daily snapshot at 02:30.

With each snapshot schedule, users can also specify the number of scheduled snapshots to be retained. Older snapshots that exceed the retention limit are automatically deleted. This user-defined limit must be equal to or lower than the global limit set by the CloudStack administrator. See [Seção 14.3, “Globally Configured Limits”](#). The limit applies only to those snapshots that are taken as part of an automatic recurring snapshot policy. Additional manual snapshots can be created and retained.

13.5.3. Incremental Snapshots and Backup

Snapshots are created on primary storage where a disk resides. After a snapshot is created, it is immediately backed up to secondary storage and removed from primary storage for optimal utilization of space on primary storage.

CloudStack does incremental backups for some hypervisors. When incremental backups are supported, every N backup is a full backup.

	VMware vSphere	Citrix XenServer	KVM
Support incremental backup	N	Y	N

13.5.4. Volume Status

When a snapshot operation is triggered by means of a recurring snapshot policy, a snapshot is skipped if a volume has remained inactive since its last snapshot was taken. A volume is considered to be inactive if it is either detached or attached to a VM that is not running. CloudStack ensures that at least one snapshot is taken since the volume last became inactive.

When a snapshot is taken manually, a snapshot is always created regardless of whether a volume has been active or not.

13.5.5. Snapshot Restore

There are two paths to restoring snapshots. Users can create a volume from the snapshot. The volume can then be mounted to a VM and files recovered as needed. Alternatively, a template may be created from the snapshot of a root disk. The user can then boot a VM from this template to effect recovery of the root disk.

Working with Usage

The Usage Server is an optional, separately-installed part of CloudStack that provides aggregated usage records which you can use to create billing integration for CloudStack. The Usage Server works by taking data from the events log and creating summary usage records that you can access using the `listUsageRecords` API call.

The usage records show the amount of resources, such as VM run time or template storage space, consumed by guest instances.

The Usage Server runs at least once per day. It can be configured to run multiple times per day.

14.1. Configuring the Usage Server

To configure the usage server:

1. Be sure the Usage Server has been installed. This requires extra steps beyond just installing the CloudStack software. See *Installing the Usage Server (Optional)* in the *Advanced Installation Guide*.
2. Faça login na interface de usuário do CloudStack como administrador.
3. Click Global Settings.
4. In Search, type usage. Find the configuration parameter that controls the behavior you want to set. See the table below for a description of the available parameters.
5. In Actions, click the Edit icon.
6. Type the desired value and click the Save icon.
7. Restart the Management Server (as usual with any global configuration change) and also the Usage Server:

```
# service cloudstack-management restart
# service cloudstack-usage restart
```

The following table shows the global configuration settings that control the behavior of the Usage Server.

Parameter Name	Descrição
<code>enable.usage.server</code>	Whether the Usage Server is active.
<code>usage.aggregation.timezone</code>	Time zone of usage records. Set this if the usage records and daily job execution are in different time zones. For example, with the following settings, the usage job will run at PST 00:15 and generate usage records for the 24 hours from 00:00:00 GMT to 23:59:59 GMT: <pre>usage.stats.job.exec.time = 00:15 usage.execution.timezone = PST usage.aggregation.timezone = GMT</pre> Valid values for the time zone are specified in Apêndice A, Time Zones

Parameter Name	Descrição
	Default: GMT
usage.execution.timezone	<p>The time zone of usage.stats.job.exec.time. Valid values for the time zone are specified in Apêndice A, Time Zones</p> <p>Default: The time zone of the management server.</p>
usage.sanity.check.interval	<p>The number of days between sanity checks. Set this in order to periodically search for records with erroneous data before issuing customer invoices. For example, this checks for VM usage records created after the VM was destroyed, and similar checks for templates, volumes, and so on. It also checks for usage times longer than the aggregation range. If any issue is found, the alert ALERT_TYPE_USAGE_SANITY_RESULT = 21 is sent.</p>
usage.stats.job.aggregation.range	<p>The time period in minutes between Usage Server processing jobs. For example, if you set it to 1440, the Usage Server will run once per day. If you set it to 600, it will run every ten hours. In general, when a Usage Server job runs, it processes all events generated since usage was last run.</p> <p>There is special handling for the case of 1440 (once per day). In this case the Usage Server does not necessarily process all records since Usage was last run. CloudStack assumes that you require processing once per day for the previous, complete day's records. For example, if the current day is October 7, then it is assumed you would like to process records for October 6, from midnight to midnight. CloudStack assumes this "midnight to midnight" is relative to the usage.execution.timezone.</p> <p>Default: 1440</p>
usage.stats.job.exec.time	<p>The time when the Usage Server processing will start. It is specified in 24-hour format (HH:MM) in the time zone of the server, which should be GMT. For example, to start the Usage job at 10:30 GMT, enter "10:30".</p> <p>If usage.stats.job.aggregation.range is also set, and its value is not 1440, then its value will be added to usage.stats.job.exec.time to get the time to run the Usage Server job again. This is repeated until 24 hours have elapsed, and the next day's processing begins again at usage.stats.job.exec.time.</p>

Parameter Name	Description
	Default: 00:15.

For example, suppose that your server is in GMT, your user population is predominantly in the East Coast of the United States, and you would like to process usage records every night at 2 AM local (EST) time. Choose these settings:

- `enable.usage.server = true`
- `usage.execution.timezone = America/New_York`
- `usage.stats.job.exec.time = 07:00`. This will run the Usage job at 2:00 AM EST. Note that this will shift by an hour as the East Coast of the U.S. enters and exits Daylight Savings Time.
- `usage.stats.job.aggregation.range = 1440`

With this configuration, the Usage job will run every night at 2 AM EST and will process records for the previous day's midnight-midnight as defined by the EST (America/New_York) time zone.



Nota

Because the special value 1440 has been used for `usage.stats.job.aggregation.range`, the Usage Server will ignore the data between midnight and 2 AM. That data will be included in the next day's run.

14.2. Setting Usage Limits

CloudStack provides several administrator control points for capping resource usage by users. Some of these limits are global configuration parameters. Others are applied at the ROOT domain and may be overridden on a per-account basis.

Aggregate limits may be set on a per-domain basis. For example, you may limit a domain and all subdomains to the creation of 100 VMs.

This section covers the following topics:

14.3. Globally Configured Limits

In a zone, the guest virtual network has a 24 bit CIDR by default. This limits the guest virtual network to 254 running instances. It can be adjusted as needed, but this must be done before any instances are created in the zone. For example, 10.1.1.0/22 would provide for ~1000 addresses.

The following table lists limits set in the Global Configuration:

Parameter Name	Definition
<code>max.account.public.ips</code>	Number of public IP addresses that can be owned by an account
<code>max.account.snapshots</code>	Number of snapshots that can exist for an account


Parameter Name	Definition
max.account.templates	Number of templates that can exist for an account
max.account.user.vms	Number of virtual machine instances that can exist for an account
max.account.volumes	Number of disk volumes that can exist for an account
max.template.iso.size	Maximum size for a downloaded template or ISO in GB
max.volume.size.gb	Maximum size for a volume in GB
network.throttling.rate	Default data transfer rate in megabits per second allowed per user (supported on XenServer)
snapshot.max.hourly	Maximum recurring hourly snapshots to be retained for a volume. If the limit is reached, early snapshots from the start of the hour are deleted so that newer ones can be saved. This limit does not apply to manual snapshots. If set to 0, recurring hourly snapshots can not be scheduled
snapshot.max.daily	Maximum recurring daily snapshots to be retained for a volume. If the limit is reached, snapshots from the start of the day are deleted so that newer ones can be saved. This limit does not apply to manual snapshots. If set to 0, recurring daily snapshots can not be scheduled
snapshot.max.weekly	Maximum recurring weekly snapshots to be retained for a volume. If the limit is reached, snapshots from the beginning of the week are deleted so that newer ones can be saved. This limit does not apply to manual snapshots. If set to 0, recurring weekly snapshots can not be scheduled
snapshot.max.monthly	Maximum recurring monthly snapshots to be retained for a volume. If the limit is reached, snapshots from the beginning of the month are deleted so that newer ones can be saved. This limit does not apply to manual snapshots. If set to 0, recurring monthly snapshots can not be scheduled.

To modify global configuration parameters, use the global configuration screen in the CloudStack UI. See [Setting Global Configuration Parameters](#)

14.4. Default Account Resource Limits

You can limit resource use by accounts. The default limits are set by using global configuration parameters, and they affect all accounts within a cloud. The relevant parameters are those beginning with `max.account`, for example: `max.account.snapshots`.


To override a default limit for a particular account, set a per-account resource limit.

1. Log in to the CloudStack UI.
2. In the left navigation tree, click Accounts.
3. Select the account you want to modify. The current limits are displayed. A value of -1 shows that there is no limit in place.
4. Click the Edit button. 

14.5. Per-Domain Limits

CloudStack allows the configuration of limits on a domain basis. With a domain limit in place, all users still have their account limits. They are additionally limited, as a group, to not exceed the resource limits set on their domain. Domain limits aggregate the usage of all accounts in the domain as well as all accounts in all subdomains of that domain. Limits set at the root domain level apply to the sum of resource usage by the accounts in all domains and sub-domains below that root domain.

To set a domain limit:

1. Log in to the CloudStack UI.
2. In the left navigation tree, click Domains.
3. Select the domain you want to modify. The current domain limits are displayed. A value of -1 shows that there is no limit in place.
4. Click the Edit button 

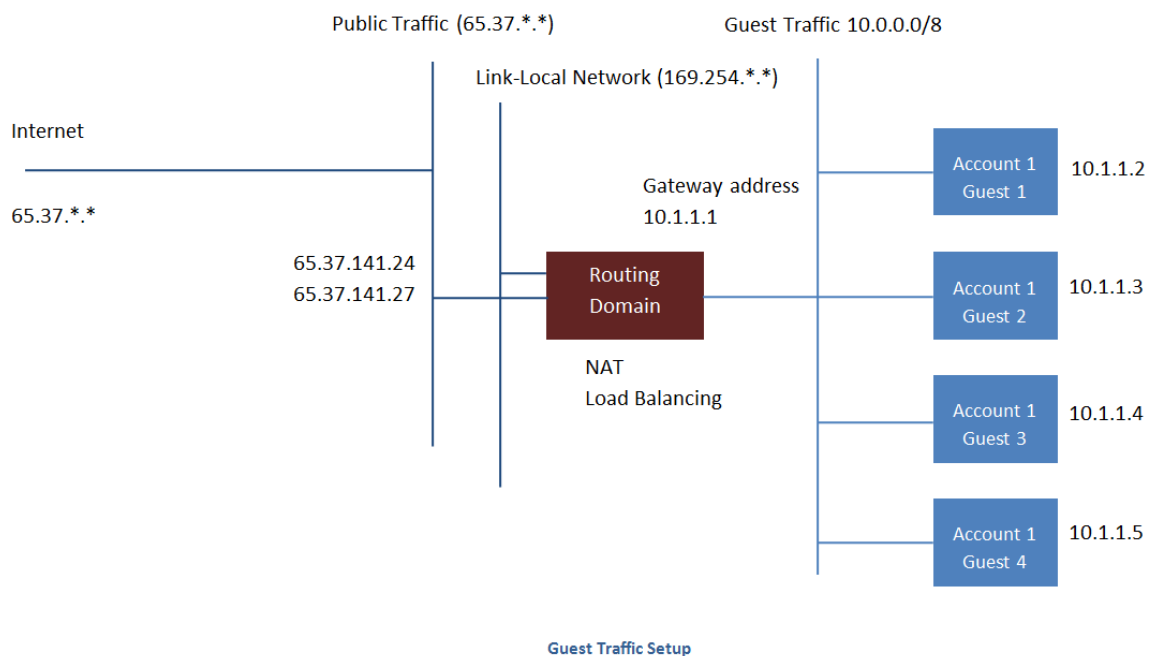
Gerenciando redes e tráfego

Em um ambiente CloudStack, máquinas virtuais hóspedes podem comunicar-se entre si usando uma infraestrutura compartilhada com segurança e com a percepção do usuário que os hóspedes têm uma rede privada. O roteador virtual do CloudStack é o principal componente a fornecer recursos de rede para o tráfego de hóspedes.

15.1. Tráfego de hóspedes

Uma rede pode transportar tráfego de hóspedes somente entre máquinas virtuais em uma zona. Máquinas virtuais em diferentes zonas não podem se comunicar usando seus endereços IP; elas devem se comunicar através de roteamento em uma rede IP pública.

This figure illustrates a typical guest traffic setup:



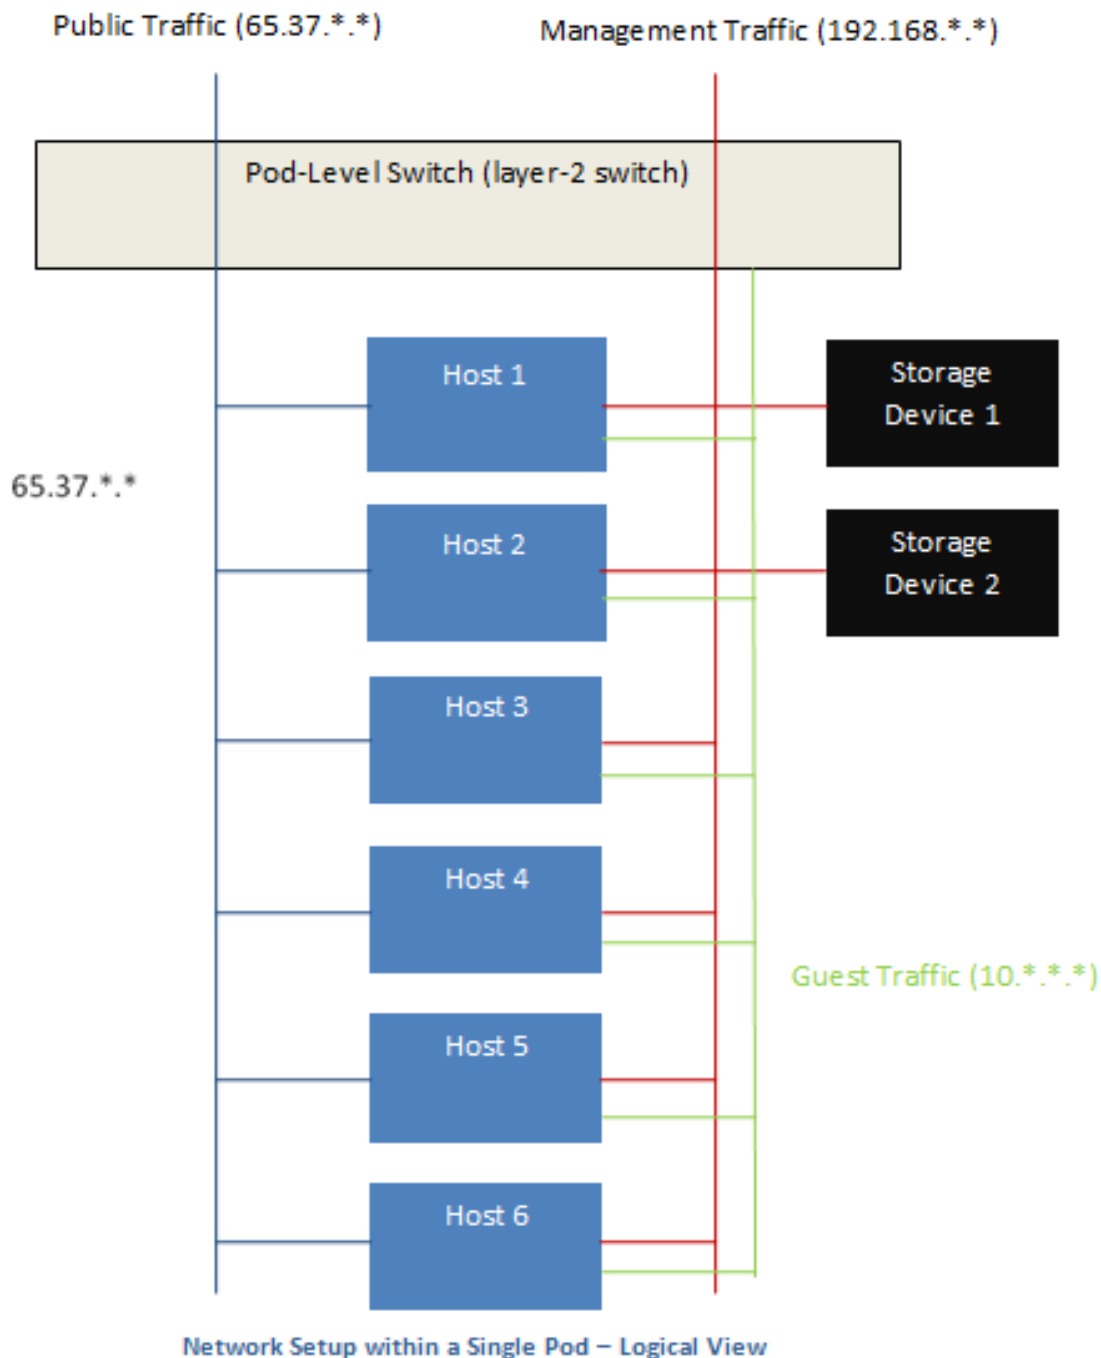
O servidor de gerenciamento automaticamente cria um roteador virtual para cada rede. Um roteador virtual é uma máquina virtual especial que executa nos hosts. Cada roteador virtual tem três interfaces de rede. Seu interface eth0 serve como o gateway para o tráfego hóspede e tem o endereço IP 10.1.1.1. Seu interface eth1 é usado pelo sistema para configurar o roteador virtual. A seu interface eth2 é assinalado um endereço IP público para tráfego público.

O roteador virtual provê DHCP e automaticamente irá assinalar um endereço IP para cada máquina virtual no intervalo de IPs assinalado para a rede. O usuário pode manualmente reconfigurar máquinas virtuais hóspedes para utilizarem diferentes endereços IP.

NAT da origem é automaticamente configurado no roteador virtual para encaminhar tráfego de saída para todas as máquinas virtuais hóspedes

15.2. Rede em um pod

The figure below illustrates network setup within a single pod. The hosts are connected to a pod-level switch. At a minimum, the hosts should have one physical uplink to each switch. Bonded NICs are supported as well. The pod-level switch is a pair of redundant gigabit switches with 10 G uplinks.



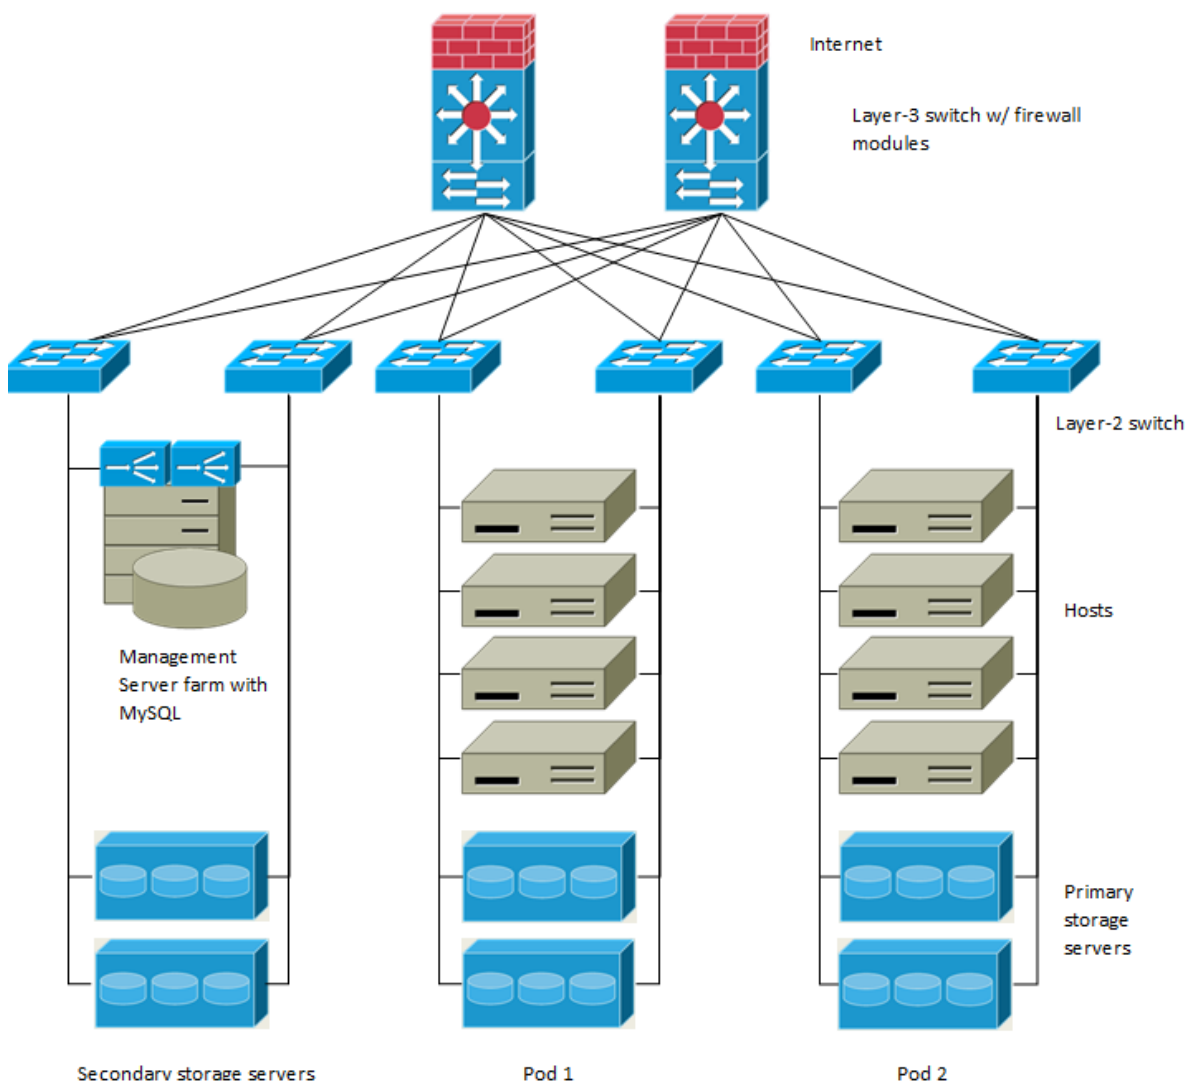
Servidores são conectados como se segue:

- Equipamentos de storage são conectados somente à rede que transporta tráfego de gerenciamento.
- Hosts são conectados a redes tanto para tráfego de gerência quanto para tráfego público.
- Hosts também são conectados a uma ou mais redes que transportam tráfego de hóspedes.

Recomendamos o uso de múltiplas placas Ethernet físicas para implementar cada interface de rede, assim como switch fabrics redundantes, para maximizar o throughput e aumentar a confiabilidade.

15.3. Rede em uma zona

The following figure illustrates the network setup within a single zone.



Um firewall para tráfego de gerência opera em modo NAT. Tipicamente à rede é assinalado um endereço IP do espaço de endereços classe B 192.168.0.0/16. A cada pod é assinalado um endereço IP no espaço de endereços classe C privado 192.168.*.0/24.

Cada zona tem seu próprio conjunto de endereços IP públicos. Endereços IP públicos de diferentes zonas não se sobrepõem.

15.4. Configuração de rede física de zona básica

Em uma rede básica, a configuração da rede física é bastante simples. Na maioria dos casos, você precisa somente configurar uma rede hóspede para transportar tráfego que é gerado pelas máquinas virtuais hóspedes. Quando você adiciona a primeira zona ao CloudStack, você configura a rede hóspede através das telas Add Zone.

15.5. Configuração de rede física de zona avançada

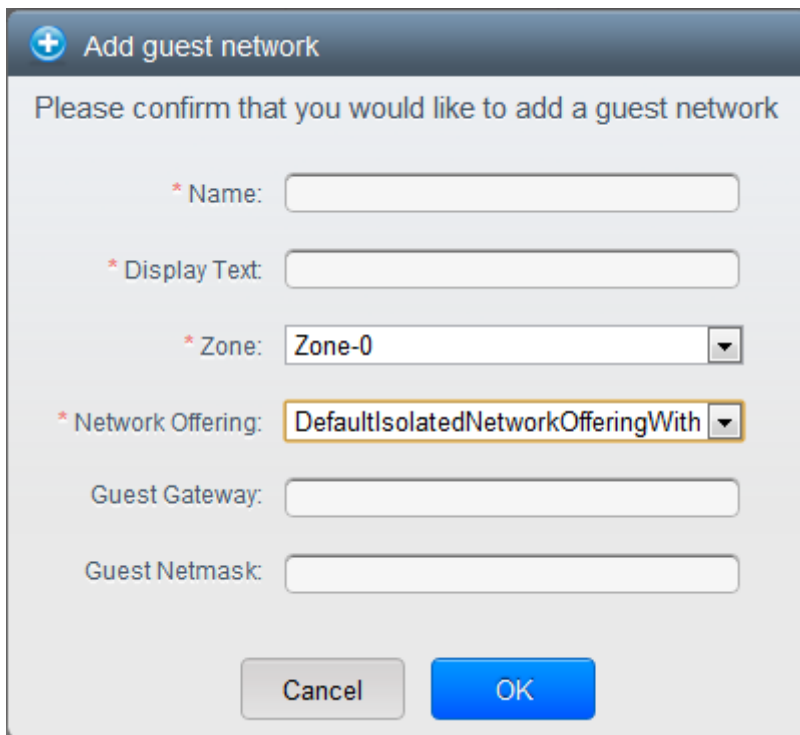
Em uma zona que usa rede avançada, você precisa informar ao servidor de gerenciamento como a rede física é configurada para transportar diferentes tipos de tráfego com isolamento.

15.5.1. Configure o tráfego hóspede na zona avançada

Nestes passos se supõe que você já está logado na interface de usuário do CloudStack. Para configurar a base da rede hóspede:

1. No painel de navegação à esquerda, selecione Infrastructure. Em Zones, clique View More, então clique a zona à qual você deseja adicionar uma rede.
2. Clique na aba Network.
3. Clique em Add guest network.

A janela Add guest network é apresentada:



4. Forneça as seguintes informações:
 - **Name.** O nome da rede. Isto será visível pelo usuário
 - **Display Text:** A descrição da rede. Isto será visível pelo usuário
 - **Zone:** A zona na qual você está configurando a rede hóspede.
 - **Network offering:** Se o administrador configurou múltiplas ofertas de rede, selecione a que você deseja usar para esta rede
 - **Guest Gateway:** O gateway que os hóspedes devem usar
 - **Guest Netmask:** A máscara de rede em uso na subnet que os hóspedes utilizarão
5. Clique em OK.

15.5.2. Configure o tráfego público na zona avançada

Em uma zona que usa rede avançada, você precisa configurar pelo menos um intervalo de endereços IP para tráfego de Internet.

15.6. Usando múltiplas redes hóspedes

Em zonas que usam rede avançada, redes adicionais para tráfego hóspede podem ser adicionadas a qualquer tempo após a instalação inicial. Você pode também customizar o nome de domínio associado com a rede especificando um sufixo DNS para cada rede.

As redes de uma máquina virtual são definidas em tempo de criação da máquina virtual. Uma máquina virtual não pode adicionar ou remover redes após ter sido criada, embora o usuário possa remover no hóspede o endereço IP da NIC de uma rede específica.

Cada máquina virtual tem somente uma rede default. A resposta do DHCP do roteador virtual irá determinar o default gateway do hóspede como aquele da rede default. Múltiplas redes non-default podem ser adicionadas a um hóspede em adição à rede default, única, requerida. O administrador pode controlar quais redes são disponíveis como rede default.

Redes adicionais podem estar disponíveis para todas as contas ou serem assinaladas a uma conta específica. Redes que estão disponíveis para todas as contas são zone-wide. Qualquer usuário com acesso à zona pode criar uma máquina virtual com acesso àquela rede. Estas redes zone-wide proveem pequeno ou nenhum isolamento entre hóspedes. Redes que são assinaladas a contas específicas proveem isolamento robusto.


15.6.1. Adicionando uma rede hóspede adicional

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Clique em Add guest network. Forneça as seguintes informações:
 - **Name.** O nome da rede. Isto será visível pelo usuário.
 - **Display Text:** A descrição da rede. Isto será visível pelo usuário.
 - **Zone.** O nome da zona a qual esta rede se aplica. Cada zona é um domínio de broadcast, e portanto cada zona tem um diferente intervalo de IP para a rede hóspede. O administrador deve configurar o intervalo de IP para cada zona.
 - **Network offering:** Se o administrador configurou múltiplas ofertas de rede, selecione a que você deseja usar para esta rede.
 - **Guest Gateway:** O gateway que os hóspedes devem usar.
 - **Guest Netmask:** A máscara de rede em uso na subnet que os hóspedes utilizarão.
4. Clique em Create.

15.6.2. Alterando a oferta de rede em uma rede hóspede

Um usuário ou administrador pode alterar a oferta de rede associada com uma rede hóspede associada.

- Faça login na interface de usuário do CloudStack como administrador ou usuário final.
- If you are changing from a network offering that uses the CloudStack virtual router to one that uses external devices as network service providers, you must first stop all the VMs on the network. See [Seção 10.6, “Stopping and Starting VMs”](#).
- Na barra de navegação à esquerda, selecione Network.

- Click the name of the network you want to modify.
- In the Details tab, click Edit. 
- Em Network Offering, escolha a nova oferta de rede, então clique em Apply.
- A prompt is displayed asking whether you want to keep the existing CIDR. This is to let you know that if you change the network offering, the CIDR will be affected. Choose No to proceed with the change.
- Wait for the update to complete. Don't try to restart VMs until the network change is complete.
- If you stopped any VMs, restart them.

15.7. Grupos de segurança

15.7.1. About Security Groups

Security groups provide a way to isolate traffic to VMs. A security group is a group of VMs that filter their incoming and outgoing traffic according to a set of rules, called ingress and egress rules. These rules filter network traffic according to the IP address that is attempting to communicate with the VM. Security groups are particularly useful in zones that use basic networking, because there is a single guest network for all guest VMs. In advanced zones, security groups are supported only on the KVM hypervisor.



Nota

In a zone that uses advanced networking, you can instead define multiple guest networks to isolate traffic to VMs.

Each CloudStack account comes with a default security group that denies all inbound traffic and allows all outbound traffic. The default security group can be modified so that all new VMs inherit some other desired set of rules.

Any CloudStack user can set up any number of additional security groups. When a new VM is launched, it is assigned to the default security group unless another user-defined security group is specified. A VM can be a member of any number of security groups. Once a VM is assigned to a security group, it remains in that group for its entire lifetime; you can not move a running VM from one security group to another.

You can modify a security group by deleting or adding any number of ingress and egress rules. When you do, the new rules apply to all VMs in the group, whether running or stopped.

If no ingress rules are specified, then no traffic will be allowed in, except for responses to any traffic that has been allowed out through an egress rule.

15.7.2. Adicionando um grupo de segurança

Um usuário ou administrador pode definir um novo grupo de segurança.

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.

2. Na barra de navegação à esquerda, selecione Network.
3. Em Select view, selecione Security Groups.
4. Clique em Add Security Group.
5. Forneça um nome e descrição.
6. Clique em OK.

O novo grupo de segurança aparece na aba Security Groups Details.

7. Para tornar útil o grupo de segurança, continue em Adicionando regras de ingresso e egresso a um grupo de segurança.

15.7.3. Security Groups in Advanced Zones (KVM Only)

CloudStack provides the ability to use security groups to provide isolation between guests on a single shared, zone-wide network in an advanced zone where KVM is the hypervisor. Using security groups in advanced zones rather than multiple VLANs allows a greater range of options for setting up guest isolation in a cloud.

Limitations

The following are not supported for this feature:

- Two IP ranges with the same VLAN and different gateway or netmask in security group-enabled shared network.
- Two IP ranges with the same VLAN and different gateway or netmask in account-specific shared networks.
- Multiple VLAN ranges in security group-enabled shared network.
- Multiple VLAN ranges in account-specific shared networks.

Security groups must be enabled in the zone in order for this feature to be used.

15.7.4. Habilitando grupos de segurança

In order for security groups to function in a zone, the security groups feature must first be enabled for the zone. The administrator can do this when creating a new zone, by selecting a network offering that includes security groups. The procedure is described in Basic Zone Configuration in the Advanced Installation Guide. The administrator can not enable security groups for an existing zone, only when creating a new zone.

15.7.5. Adicionando regras de ingresso e egresso a um grupo de segurança

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Em Select view, selecione Security Groups, então clique no grupo de segurança que você deseja.
4. Para adicionar uma regra de ingresso, clique na aba Ingress Rules e preencha os campos seguintes para especificar qual tráfego de rede é permitido em instâncias de máquinas virtuais

neste grupo de segurança. Se nenhuma regra de ingresso é especificada, então nenhum tráfego será permitido entrar, exceto respostas a qualquer tráfego que tenha sido permitido sair por uma regra de egresso.

- **Add by CIDR/Account.** Indica se a fonte do tráfego será definida pelo endereço IP (CIDR) ou por um grupo de segurança existente em uma conta do CloudStack (Account). Selecione Account se você deseja permitir tráfego de entrada oriundo de todas as máquinas virtuais em outro grupo de segurança
- **Protocol.** O protocolo de rede que as fontes utilizarão para enviar tráfego ao grupo de segurança. TCP e UDP são usados tipicamente para intercâmbio de dados e comunicações de usuários finais. ICMP é usado tipicamente para enviar mensagens de erro e dados de monitoração de rede.
- **Start Port, End Port.** (somente para TCP e UDP) Um intervalo de listening ports que são os destinos do tráfego entrante. Se você está abrindo uma única porta, use o mesmo número em ambos os campos.
- **ICMP Type, ICMP Code.** (somente ICMP) O tipo de mensagem e o código de erro que serão aceitos.
- **CIDR.** (Adição somente por CIDR) Para aceitar tráfego somente de endereços IP em um bloco de endereços específico, informe um CIDR ou uma lista de CIDRs separados por vírgulas. O CIDR é o endereço IP base do tráfego entrante. Por exemplo, 192.168.0.0/22. Para permitir todos os CIDRs, especifique 0.0.0.0/0.
- **Account, Security Group.** (Adição somente por conta) Para aceitar somente tráfego de outro grupo de segurança, informe a conta do CloudStack e nome de um grupo de segurança que já esteja definido naquela conta. Para permitir tráfego entre máquinas virtuais no grupo de segurança que você está editando, informe o mesmo nome que você usou no passo 7.

O exemplo a seguir permite acesso HTTP entrante de qualquer lugar:

Protocol	Start Port	End Port	CIDR	Add
TCP	80	80	0.0.0.0/0	Add

5. Para adicionar uma regra de egresso, clique na aba Egress Rules e preencha os campos seguintes para especificar qual tipo de tráfego de rede é permitido ser enviado de instâncias de máquinas virtuais neste grupo de segurança. Se nenhuma regra de egresso é especificada, então todo tráfego será permitido sair. Uma vez que regras de egresso são especificadas, os seguintes tipos de tráfego são permitidos sair: tráfego especificado regras de egresso; queries a servidores DNS e DHCP; e respostas a qualquer tráfego que tenha sido permitido entrar através de uma regra de ingresso
- **Add by CIDR/Account.** Indica se o destino do tráfego será definido pelo endereço IP (CIDR) ou por um grupo de segurança existente em uma conta do CloudStack (Account). Selecione

Account se você deseja permitir tráfego de saída destinado a todas as máquinas virtuais em outro grupo de segurança.

- **Protocol.** O protocolo de rede que máquinas virtuais usarão para enviar tráfego de saída. TCP e UDP são usados tipicamente para intercâmbio de dados e comunicações de usuários finais. ICMP é usado tipicamente para enviar mensagens de erro e dados de monitoração de rede.
- **Start Port, End Port.** (somente para TCP e UDP) Um intervalo de listening ports que são os destinos do tráfego sainte. Se você está abrindo uma única porta, use o mesmo número em ambos os campos.
- **ICMP Type, ICMP Code.** (somente ICMP) O tipo de mensagem e o código de erro que serão enviados.
- **CIDR.** (Adição somente por CIDR) Para enviar tráfego somente a endereços IP em um bloco de endereços específico, informe um CIDR ou uma lista de CIDRs separados por vírgulas. O CIDR é o endereço IP base do tráfego entrante. Por exemplo, 192.168.0.0/22. Para permitir todos os CIDRs, especifique 0.0.0.0/0.
- **Account, Security Group.** (Adição somente por conta) Para permitir o envio de tráfego a outro grupo de segurança, informe a conta do CloudStack e nome de um grupo de segurança que já esteja definido naquela conta. Para permitir tráfego entre máquinas virtuais no grupo de segurança que você está editando, informe o nome do grupo.

6. Clique em Add.

15.8. Firewalls e balanceadores de carga externos

O CloudStack é capaz de substituir seu roteador virtual por um equipamento Juniper SRX externo e um balanceador de carga externo NetScaler ou F5 para serviços de gateway e balanceamento de carga. Neste caso, as máquinas virtuais usam o SRX como seu gateway.

15.8.1. Sobre a utilização do balanceador de carga NetScaler

O Citrix NetScaler é suportado como um elemento externo de rede para balanceamento de carga em zonas que usam rede avançada (também chamadas zonas avançadas). Configure um balanceador de carga externo quando você quiser prover balanceamento de carga por outros meios que o roteador virtual provido pelo CloudStack.

O NetScaler pode ser configurado em modo direto (fora do firewall). Ele deve ser adicionado antes que qualquer regra de balanceamento de carga seja implementada em máquinas virtuais da zona.

O comportamento funcional do NetScaler com o CloudStack é o mesmo que é descrito na documentação do CloudStack para uso de um balanceador de carga externo F5. A única exceção é que o F5 suporta domínios de roteamento, e o NetScaler não. O NetScaler ainda não pode ser usado como um firewall.

O Citrix NetScaler é oferecido em três variantes. A tabela a seguir sumariza como estas variantes são tratadas no CloudStack.

Tipo de NetScaler ADC	Descrição das capacidades	Recursos suportados no CloudStack
MPX	Dispositivo físico. Capaz de inspeção profunda de pacotes. Pode atuar como application firewall e balanceador de carga	Em zonas avançadas, a funcionalidade de balanceador de carga é suportada sem limitações. Em zonas básicas,

Tipo de NetScaler ADC	Descrição das capacidades	Recursos suportados no CloudStack
		NAT estática, elastic IP (EIP), e elastic load balancing (ELB) também são providos
VPX	Dispositivo virtual. Pode ser executado como máquina virtual nos hipervisores XenServer, ESXi, e Hyper-V. Mesma funcionalidade que MPX	Suportado somente no ESXi. Mesmo suporte funcional que MPX. O CloudStack tratará VPX e MPX como o mesmo tipo de dispositivo
SDX	Dispositivo físico. Pode criar múltiplas instâncias VPX completamente isoladas em um único dispositivo para suportar uso por múltiplos hóspedes	O CloudStack irá dinamicamente aprovisionar, configurar, e gerenciar o ciclo de vida de instâncias VPX no SDX. Instâncias fornecidas são automaticamente adicionadas ao CloudStack – nenhuma configuração manual pelo administrador é requerida. Uma vez que uma instância VPX é adicionada ao CloudStack, ela é tratada da mesma forma que uma VPX em um host ESXi.

15.8.2. Configuring SNMP Community String on a RHEL Server


The SNMP Community string is similar to a user id or password that provides access to a network device, such as router. This string is sent along with all SNMP requests. If the community string is correct, the device responds with the requested information. If the community string is incorrect, the device discards the request and does not respond.

The NetScaler device uses SNMP to communicate with the VMs. You must install SNMP and configure SNMP Community string for a secure communication between the NetScaler device and the RHEL machine.

1. Ensure that you installed SNMP on RedHat. If not, run the following command:

```
yum install net-snmp-utils
```

2. Edit the /etc/snmp/snmpd.conf file to allow the SNMP polling from the NetScaler device.
 - a. Map the community name into a security name (local and mynetwork, depending on where the request is coming from):



Nota

Use a strong password instead of public when you edit the following table.

```
#      sec.name  source      community
com2sec  local      localhost   public
com2sec  mynetwork  0.0.0.0     public
```



Nota

Setting to 0.0.0.0 allows all IPs to poll the NetScaler server.

- b. Map the security names into group names:

```
#      group.name  sec.model  sec.name
group  MyRWGroup     v1         local
group  MyRWGroup     v2c        local
group  MyROGroup     v1         mynetwork
group  MyROGroup     v2c        mynetwork
```

- c. Create a view to allow the groups to have the permission to:

```
incl/excl subtree mask view all included .1
```

- d. Grant access with different write permissions to the two groups to the view you created.

```
# context      sec.model  sec.level  prefix  read  write  notif
access        MyROGroup ""         any noauth  exact  all   none  none
access        MyRWGroup ""         any noauth  exact  all   all   all
```

3. Unblock SNMP in iptables.

```
iptables -A INPUT -p udp --dport 161 -j ACCEPT
```

4. Start the SNMP service:

```
service snmpd start
```

5. Ensure that the SNMP service is started automatically during the system startup:

```
chkconfig snmpd on
```

15.8.3. Configuração inicial de firewalls e balanceadores de carga externos

Quando a primeira máquina virtual é criada para uma nova conta, o CloudStack programa o firewall e o balanceador de carga externos para trabalhar com a máquina virtual. Os seguintes objetos são criados no firewall:

- Uma nova interface lógica para conectar à VLAN privada da conta. O IP da interface é sempre o primeiro IP da subnet privada da conta (e.g. 10.1.1.1).

- Uma regra NAT de origem que encaminha todo o tráfego sainte da VLAN privada da conta para a Internet pública, usando o endereço IP público da conta como o endereço de origem
- Um contador de filtro de firewall que registra o número de bytes do tráfego sainte da conta

Os seguintes objetos são criados no balanceador de carga:

- Uma nova VLAN que corresponde à VLAN na zona fornecida para a conta
- Um IP para a VLAN. Este é sempre o segundo IP da subnet privada da conta (e.g. 10.1.1.2).

15.8.4. Configuração continuada de firewalls e balanceadores de carga externos

Ações adicionais de usuários (e.g. configuração de um encaminhamento de porta) causará programação adicional do firewall e balanceador de carga. Um usuário pode requisitar endereços IP públicos adicionais e encaminhamento de tráfego recebido nestes IPs para máquinas virtuais específicas. Isto é executado através da habilitação de NAT estático para um endereço IP público, assinalando o IP a uma máquina virtual, e especificando um conjunto de protocolos e intervalos de portas a liberar. Quando uma regra NAT estática é criada, o CloudStack programa o firewall externo da zona com os seguintes objetos:

- Uma regra NAT estática que mapeia o endereço IP público ao endereço IP privado de uma máquina virtual.
- Uma política de segurança que permite tráfego no conjunto de protocolos e intervalos de porta que são especificados.
- Um contador de filtro de firewall que registra o número de bytes do tráfego entrante no IP público.

O número de bytes entrantes e saintes através de NAT de origem, NAT estático, e regras de balanceamento de carga é medido e salvo em cada elemento externo. Este dado é coletado regularmente e armazenado no database do CloudStack.

15.8.5. Configuring AutoScale

AutoScaling allows you to scale your back-end services or application VMs up or down seamlessly and automatically according to the conditions you define. With AutoScaling enabled, you can ensure that the number of VMs you are using seamlessly scale up when demand increases, and automatically decreases when demand subsides. Using AutoScaling, you can automatically shut down instances you don't need, or launch new instances, depending on demand.

NetScaler AutoScaling is designed to seamlessly launch or terminate VMs based on user-defined conditions. Conditions for triggering a scaleup or scaledown action can vary from a simple use case like monitoring the CPU usage of a server to a complex use case of monitoring a combination of server's responsiveness and its CPU usage. For example, you can configure AutoScaling to launch an additional VM whenever CPU usage exceeds 80 percent for 15 minutes, or to remove a VM whenever CPU usage is less than 20 percent for 30 minutes.

CloudStack uses the NetScaler load balancer to monitor all aspects of a system's health and work in unison with CloudStack to initiate scale-up or scale-down actions.



Nota

AutoScale is supported on NetScaler Release 10 Build 73.e and beyond.

Prerequisites

Before you configure an AutoScale rule, consider the following:

- Ensure that the necessary template is prepared before configuring AutoScale. When a VM is deployed by using a template and when it comes up, the application should be up and running.



Nota

If the application is not running, the NetScaler device considers the VM as ineffective and continues provisioning the VMs unconditionally until the resource limit is exhausted.

- Deploy the templates you prepared. Ensure that the applications come up on the first boot and is ready to take the traffic. Observe the time requires to deploy the template. Consider this time when you specify the quiet time while configuring AutoScale.
- The AutoScale feature supports the SNMP counters that can be used to define conditions for taking scale up or scale down actions. To monitor the SNMP-based counter, ensure that the SNMP agent is installed in the template used for creating the AutoScale VMs, and the SNMP operations work with the configured SNMP community and port by using standard SNMP managers. For example, see [Seção 15.8.2, “Configuring SNMP Community String on a RHEL Server”](#) to configure SNMP on a RHEL machine.
- Ensure that the `endpoint.url` parameter present in the Global Settings is set to the Management Server API URL. For example, `http://10.102.102.22:8080/client/api`. In a multi-node Management Server deployment, use the virtual IP address configured in the load balancer for the management server’s cluster. Additionally, ensure that the NetScaler device has access to this IP address to provide AutoScale support.

If you update the `endpoint.url`, disable the AutoScale functionality of the load balancer rules in the system, then enable them back to reflect the changes. For more information see [Updating an AutoScale Configuration](#)

- If the API Key and Secret Key are regenerated for an AutoScale user, ensure that the AutoScale functionality of the load balancers that the user participates in are disabled and then enabled to reflect the configuration changes in the NetScaler.
- In an advanced Zone, ensure that at least one VM should be present before configuring a load balancer rule with AutoScale. Having one VM in the network ensures that the network is in implemented state for configuring AutoScale.

Configuração

Especifique o seguinte:

AutoScale Configuration Wizard

Template: RHEL62

Compute offering: Small Instance

* Min Instances: 1 * Max Instances: 4

Scale Up Policy

* Duration(in sec): 60

Counter	Operator	Threshold	Add
Linux User CPU - percentage	greater-than		Add
Response Time - microseconds	greater-than	1000	X

Scale Down Policy

* Duration(in sec): 60

Counter	Operator	Threshold	Add

Cancel Apply

- **Template:** A template consists of a base OS image and application. A template is used to provision the new instance of an application on a scaleup action. When a VM is deployed from a template, the VM can start taking the traffic from the load balancer without any admin intervention. For example, if the VM is deployed for a Web service, it should have the Web server running, the database connected, and so on.
- **Compute offering:** A predefined set of virtual hardware attributes, including CPU speed, number of CPUs, and RAM size, that the user can select when creating a new virtual machine instance. Choose one of the compute offerings to be used while provisioning a VM instance as part of scaleup action.
- **Min Instance:** The minimum number of active VM instances that is assigned to a load balancing rule. The active VM instances are the application instances that are up and serving the traffic, and are being load balanced. This parameter ensures that a load balancing rule has at least the configured number of active VM instances available to serve the traffic.

**Nota**

If an application, such as SAP, running on a VM instance is down for some reason, the VM is then not counted as part of Min Instance parameter, and the AutoScale feature initiates a scaleup action if the number of active VM instances is below the configured value. Similarly, when an application instance comes up from its earlier down state, this application instance is counted as part of the active instance count and the AutoScale process initiates a scaledown action when the active instance count breaches the Max instance value.

- **Max Instance:** Maximum number of active VM instances that **should be assigned to** a load balancing rule. This parameter defines the upper limit of active VM instances that can be assigned to a load balancing rule.

Specifying a large value for the maximum instance parameter might result in provisioning large number of VM instances, which in turn leads to a single load balancing rule exhausting the VM instances limit specified at the account or domain level.

**Nota**

If an application, such as SAP, running on a VM instance is down for some reason, the VM is not counted as part of Max Instance parameter. So there may be scenarios where the number of VMs provisioned for a scaleup action might be more than the configured Max Instance value. Once the application instances in the VMs are up from an earlier down state, the AutoScale feature starts aligning to the configured Max Instance value.

Specify the following scale-up and scale-down policies:

- **Duration:** The duration, in seconds, for which the conditions you specify must be true to trigger a scaleup action. The conditions defined should hold true for the entire duration you specify for an AutoScale action to be invoked.
- **Counter:** The performance counters expose the state of the monitored instances. By default, CloudStack offers four performance counters: Three SNMP counters and one NetScaler counter. The SNMP counters are Linux User CPU, Linux System CPU, and Linux CPU Idle. The NetScaler counter is ResponseTime. The root administrator can add additional counters into CloudStack by using the CloudStack API.
- **Operator:** The following five relational operators are supported in AutoScale feature: Greater than, Less than, Less than or equal to, Greater than or equal to, and Equal to.
- **Threshold:** Threshold value to be used for the counter. Once the counter defined above breaches the threshold value, the AutoScale feature initiates a scaleup or scaledown action.
- **Add:** Click Add to add the condition.

Additionally, if you want to configure the advanced settings, click Show advanced settings, and specify the following:


- **Polling interval:** Frequency in which the conditions, combination of counter, operator and threshold, are to be evaluated before taking a scale up or down action. The default polling interval is 30 seconds.
- **Quiet Time:** This is the cool down period after an AutoScale action is initiated. The time includes the time taken to complete provisioning a VM instance from its template and the time taken by an application to be ready to serve traffic. This quiet time allows the fleet to come up to a stable state before any action can take place. The default is 300 seconds.
- **Destroy VM Grace Period:** The duration in seconds, after a scaledown action is initiated, to wait before the VM is destroyed as part of scaledown action. This is to ensure graceful close of any pending sessions or transactions being served by the VM marked for destroy. The default is 120 seconds.
- **Security Groups:** Security groups provide a way to isolate traffic to the VM instances. A security group is a group of VMs that filter their incoming and outgoing traffic according to a set of rules, called ingress and egress rules. These rules filter network traffic according to the IP address that is attempting to communicate with the VM.
- **Disk Offerings:** A predefined set of disk size for primary data storage.
- **SNMP Community:** The SNMP community string to be used by the NetScaler device to query the configured counter value from the provisioned VM instances. Default is public.
- **SNMP Port:** The port number on which the SNMP agent that run on the provisioned VMs is listening. Default port is 161.
- **User:** This is the user that the NetScaler device use to invoke scaleup and scaledown API calls to the cloud. If no option is specified, the user who configures AutoScaling is applied. Specify another user name to override.
- **Apply:** Click Apply to create the AutoScale configuration.

Disabling and Enabling an AutoScale Configuration

If you want to perform any maintenance operation on the AutoScale VM instances, disable the AutoScale configuration. When the AutoScale configuration is disabled, no scaleup or scaledown action is performed. You can use this downtime for the maintenance activities. To disable the

AutoScale configuration, click the Disable AutoScale  button.

The button toggles between enable and disable, depending on whether AutoScale is currently enabled or not. After the maintenance operations are done, you can enable the AutoScale configuration back.

To enable, open the AutoScale configuration page again, then click the Enable AutoScale  button.

Updating an AutoScale Configuration

You can update the various parameters and add or delete the conditions in a scaleup or scaledown rule. Before you update an AutoScale configuration, ensure that you disable the AutoScale load balancer rule by clicking the Disable AutoScale button.

After you modify the required AutoScale parameters, click Apply. To apply the new AutoScale policies, open the AutoScale configuration page again, then click the Enable AutoScale button.

Runtime Considerations

- An administrator should not assign a VM to a load balancing rule which is configured for AutoScale.
- Before a VM provisioning is completed if NetScaler is shutdown or restarted, the provisioned VM cannot be a part of the load balancing rule though the intent was to assign it to a load balancing rule. To workaround, rename the AutoScale provisioned VMs based on the rule name or ID so at any point of time the VMs can be reconciled to its load balancing rule.
- Making API calls outside the context of AutoScale, such as destroyVM, on an autoscaled VM leaves the load balancing configuration in an inconsistent state. Though VM is destroyed from the load balancer rule, NetScaler continues to show the VM as a service assigned to a rule.

15.9. Regras de balanceamento de carga

Um usuário ou administrador do CloudStack pode criar regras de balanceamento de carga que distribuem o tráfego recebido em um endereço IP público por uma ou mais máquinas virtuais. Um usuário cria uma regra, especifica um algoritmo e assinala a regra a um conjunto de máquinas virtuais.



Nota

Se você cria regras de balanceamento de carga enquanto usando um oferta de serviço de rede que inclui um equipamento externo de balanceamento de carga, como o NetScaler, e depois altera a oferta de serviço para um que usa o roteador virtual do CloudStack, você deve criar uma regra no firewall do roteador virtual para cada uma das regras de balanceamento de carga existentes, de forma que elas possam continuar funcionando.

15.9.1. Adding a Load Balancer Rule

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Click the name of the network where you want to load balance the traffic.
4. Clique em View IP Addresses.
5. Clique no endereço IP para o qual você deseja criar a regra, então clique na aba Configuration.
6. No nó Load Balancing do diagrama, clique em View All.

In a Basic zone, you can also create a load balancing rule without acquiring or selecting an IP address. CloudStack internally assign an IP when you create the load balancing rule, which is listed in the IP Addresses page when the rule is created.

To do that, select the name of the network, then click Add Load Balancer tab. Continue with [7](#).

7. Fill in the following:

- **Name:** Um nome para a regra de balanceamento de carga.
 - **Public Port:** The port receiving incoming traffic to be balanced.
 - **Private Port:** A porta que as máquinas virtuais usarão para receber o tráfego.
 - **Algorithm:** Choose the load balancing algorithm you want CloudStack to use. CloudStack supports a variety of well-known algorithms. If you are not familiar with these choices, you will find plenty of information about them on the Internet.
 - **Stickiness:** (Optional) Click Configure and choose the algorithm for the stickiness policy. See Sticky Session Policies for Load Balancer Rules.
 - **AutoScale:** Click Configure and complete the AutoScale configuration as explained in [Seção 15.8.5, “Configuring AutoScale”](#).
8. Click Add VMs, then select two or more VMs that will divide the load of incoming traffic, and click Apply.

The new load balancer rule appears in the list. You can repeat these steps to add more load balancer rules for this IP address.

15.9.2. Sticky Session Policies for Load Balancer Rules

Sticky sessions are used in Web-based applications to ensure continued availability of information across the multiple requests in a user's session. For example, if a shopper is filling a cart, you need to remember what has been purchased so far. The concept of "stickiness" is also referred to as persistence or maintaining state.

Any load balancer rule defined in CloudStack can have a stickiness policy. The policy consists of a name, stickiness method, and parameters. The parameters are name-value pairs or flags, which are defined by the load balancer vendor. The stickiness method could be load balancer-generated cookie, application-generated cookie, or source-based. In the source-based method, the source IP address is used to identify the user and locate the user's stored data. In the other methods, cookies are used. The cookie generated by the load balancer or application is included in request and response URLs to create persistence. The cookie name can be specified by the administrator or automatically generated. A variety of options are provided to control the exact behavior of cookies, such as how they are generated and whether they are cached.

For the most up to date list of available stickiness methods, see the CloudStack UI or call `listNetworks` and check the `SupportedStickinessMethods` capability.

15.10. Guest IP Ranges

The IP ranges for guest network traffic are set on a per-account basis by the user. This allows the users to configure their network in a fashion that will enable VPN linking between their guest network and their clients.

15.11. Obtendo um novo endereço IP

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Clique no nome da rede com a qual você quer trabalhar.

4. Clique em View IP Addresses.
5. Clique em Acquire New IP, e clique Yes no diálogo de confirmação.

Você deve confirmar porque, tipicamente, endereços IP são um recurso limitado. Em alguns momentos, o novo endereço IP deve aparecer no estado Allocated. Agora você pode usar o endereço IP no encaminhamento de porta ou regras NAT estáticas.

15.12. Liberando um endereço IP

When the last rule for an IP address is removed, you can release that IP address. The IP address still belongs to the VPC; however, it can be picked up for any guest network again.

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Clique no nome da rede com a qual você quer trabalhar.
4. Clique em View IP Addresses.
5. Clique no endereço IP que você deseja liberar.

6.  Click the Release IP button.

15.13. NAT estática

A static NAT rule maps a public IP address to the private IP address of a VM in order to allow Internet traffic into the VM. The public IP address always remains the same, which is why it is called “static” NAT. This section tells how to enable or disable static NAT for a particular IP address.

15.13.1. Habilitando ou desabilitando NAT estática

Se regras de encaminhamento de portas já estão em efeito para um endereço IP, você não pode habilitar NAT estática para este IP.

Se uma máquina virtual hóspede faz parte de mais de uma rede, regras de NAT estática funcionarão somente se elas estão definidas na rede default.

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Clique no nome da rede com a qual você quer trabalhar.
4. Clique em View IP Addresses.
5. Clique no endereço IP com o qual você deseja trabalhar.

6.  Click the Static NAT button.

The button toggles between Enable and Disable, depending on whether static NAT is currently enabled for the IP address.

7. If you are enabling static NAT, a dialog appears where you can choose the destination VM and click Apply.

15.14. Encaminhamento de IP e firewall

By default, all incoming traffic to the public IP address is rejected. All outgoing traffic from the guests is also blocked by default.

To allow outgoing traffic, follow the procedure in [Seção 15.14.1, “Creating Egress Firewall Rules in an Advanced Zone”](#).

To allow incoming traffic, users may set up firewall rules and/or port forwarding rules. For example, you can use a firewall rule to open a range of ports on the public IP address, such as 33 through 44. Then use port forwarding rules to direct traffic from individual ports within that range to specific ports on user VMs. For example, one port forwarding rule could route incoming traffic on the public IP's port 33 to port 100 on one user VM's private IP. For more information, see [Seção 15.14.2, “Regras de firewall”](#) and [Seção 15.14.3, “Encaminhamento de Porta”](#).

15.14.1. Creating Egress Firewall Rules in an Advanced Zone



Nota

The egress firewall rules are supported only on virtual routers.

The egress traffic originates from a private network to a public network, such as the Internet. By default, the egress traffic is blocked, so no outgoing traffic is allowed from a guest network to the Internet. However, you can control the egress traffic in an Advanced zone by creating egress firewall rules. When an egress firewall rule is applied, the traffic specific to the rule is allowed and the remaining traffic is blocked. When all the firewall rules are removed the default policy, Block, is applied.

Consider the following scenarios to apply egress firewall rules:

- Allow the egress traffic from specified source CIDR. The Source CIDR is part of guest network CIDR.
- Allow the egress traffic with destination protocol TCP,UDP,ICMP, or ALL.
- Allow the egress traffic with destination protocol and port range. The port range is specified for TCP, UDP or for ICMP type and code.

To configure an egress firewall rule:

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. In Select view, choose Guest networks, then click the Guest network you want.
4. To add an egress rule, click the Egress rules tab and fill out the following fields to specify what type of traffic is allowed to be sent out of VM instances in this guest network:

CIDR	Protocol	Start Port	End Port	Add
<input type="text"/>	TCP	<input type="text"/>	<input type="text"/>	<input type="button" value="Add"/>
10.1.1.0/24	TCP	22	22	<input type="button" value="✕"/>

- **CIDR:** (Add by CIDR only) To send traffic only to the IP addresses within a particular address block, enter a CIDR or a comma-separated list of CIDRs. The CIDR is the base IP address of the destination. For example, 192.168.0.0/22. To allow all CIDRs, set to 0.0.0.0/0.
- **Protocol:** The networking protocol that VMs uses to send outgoing traffic. The TCP and UDP protocols are typically used for data exchange and end-user communications. The ICMP protocol is typically used to send error messages or network monitoring data.
- **Start Port, End Port:** (TCP, UDP only) A range of listening ports that are the destination for the outgoing traffic. If you are opening a single port, use the same number in both fields.
- **ICMP Type, ICMP Code:** (ICMP only) The type of message and error code that are sent.

5. Clique em Add.

15.14.2. Regras de firewall

Por default, todo o tráfego entrante no endereço IP público é rejeitado pelo firewall. Para permitir tráfego externo, você pode abrir portas no firewall especificando regras de firewall. Opcionalmente, você pode especificar um ou mais CIDRs para filtrar os IPs de origem. Isto é útil quando você deseja permitir tráfego entrante somente de certos endereços IP.

Você não pode usar regras de firewall para abrir portas para um endereço IP elástico. Quando um IP elástico é usado, acesso externo é controlado pelo uso de grupos de segurança. Veja [Seção 15.7.2, "Adicionando um grupo de segurança"](#).

In an advanced zone, you can also create egress firewall rules by using the virtual router. For more information, see [Seção 15.14.1, "Creating Egress Firewall Rules in an Advanced Zone"](#).

Regras de firewall podem ser criadas usando a aba Firewall no interface de usuário do Servidor de gerenciamento. Por default, esta aba não é apresentada quando o CloudStack é instalado. Para exibir a aba Firewall, o administrador do CloudStack deve configurar o parâmetro global `firewall.rule.ui.enabled` como "true."

Para criar uma regra de firewall:

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Clique no nome da rede com a qual você quer trabalhar.
4. Clique em View IP Addresses.
5. Clique no endereço IP com o qual você deseja trabalhar.
6. Clique na aba Configuration e preencha os seguintes valores.

- **Source CIDR.** (Opcional) Para aceitar tráfego somente de endereços IP em um bloco de endereços específico, informe um CIDR ou uma lista de CIDRs separados por vírgulas.. Exemplo: 192.168.0.0/22. Deixe vazio para permitir todos os CIDRs.
- **Protocol.** O protocolo de comunicação em uso na(s) porta(s) aberta(s).
- **Start Port e End Port.** A(s) porta(s) que você deseja abrir no firewall. Se você está abrindo um única porta, use o mesmo número em ambos os campos
- **ICMP Type e ICMP Code.** Usados somente se Protocol é configurado como ICMP. Proveem o tipo e o código requeridos pelo protocolo ICMP para preencher o cabeçalho ICMP. Consulte a documentação do ICMP para mais detalhes se você não tem certeza do que informar

7. Clique em Add.

15.14.3. Encaminhamento de Porta

A port forward service is a set of port forwarding rules that define a policy. A port forward service is then applied to one or more guest VMs. The guest VM then has its inbound network access managed according to the policy defined by the port forwarding service. You can optionally specify one or more CIDRs to filter the source IPs. This is useful when you want to allow only incoming requests from certain IP addresses to be forwarded.

A guest VM can be in any number of port forward services. Port forward services can be defined but have no members. If a guest VM is part of more than one network, port forwarding rules will function only if they are defined on the default network

You cannot use port forwarding to open ports for an elastic IP address. When elastic IP is used, outside access is instead controlled through the use of security groups. See Security Groups.

To set up port forwarding:

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. If you have not already done so, add a public IP address range to a zone in CloudStack. See Adding a Zone and Pod in the Installation Guide.
3. Add one or more VM instances to CloudStack.
4. In the left navigation bar, click Network.
5. Click the name of the guest network where the VMs are running.
6. Choose an existing IP address or acquire a new IP address. See [Seção 15.11, “Obtendo um novo endereço IP”](#). Click the name of the IP address in the list.
7. Click the Configuration tab.
8. No nó Port Forwarding do diagrama, clique em View All.
9. Fill in the following:
 - **Public Port.** The port to which public traffic will be addressed on the IP address you acquired in the previous step.
 - **Private Port.** The port on which the instance is listening for forwarded public traffic.
 - **Protocol.** The communication protocol in use between the two ports

10. Clique em Add.

15.15. Balanceamento de carga de IP

O usuário pode escolher associar o mesmo IP público para vários hóspedes. O CloudStack implementa um balanceador de carga em nível de TCP com as seguintes políticas.

- Round-robin
- Menor quantidade de conexões
- IP de origem

Isto é similar a encaminhamento de porta, mas o destino pode ser múltiplos endereços IP.

15.16. DNS e DHCP

O roteador virtual provê serviços DNS e DHCP aos hóspedes. Ele faz proxy de requisições DNS ao servidor DNS configurado na Availability Zone.

15.17. VPN

Donos de contas no CloudStack podem criar redes privadas virtuais (VPN) para acessar suas máquinas virtuais. Se a rede hóspede é instanciada de uma oferta de rede que oferece o serviço de acesso remoto por VPN, o roteador virtual (baseado na máquina virtual de sistema) é usado para prover o serviço. O CloudStack provê um serviço de acesso remoto por VPN baseado em L2TP-over-IPsec para redes virtuais hóspedes. Como cada rede obtém seu próprio roteador virtual, as VPNs não são compartilhadas pelas redes. Clientes VPN nativos em Windows, Mac OS X e iOS podem ser usados para conectar às redes hóspedes. O dono da conta pode criar e gerenciar usuários para suas VPNs. O CloudStack não usa seu database de contas com este propósito, mas usa uma tabela distinta. O database de usuários de VPN é compartilhado por todas as VPNs criadas pelo dono da conta. Todos os usuários de VPN obtêm acesso a todas as VPNs criadas pelo dono da conta.



Nota

Certifique-se de que nem todo tráfego passa pela VPN. Isto é, a rota estabelecida pela VPN deve ser apenas para a rede hóspede e não para todo o tráfego.


- **Road Warrior / Remote Access.** Usuários desejam ser capazes de se conectar com segurança de casa ou do escritório a uma rede privada na nuvem. Tipicamente, o endereço IP do cliente que se conecta é dinâmico e não pode ser pré-configurado no servidor VPN.
- **Site to Site.** In this scenario, two private subnets are connected over the public Internet with a secure VPN tunnel. The cloud user's subnet (for example, an office network) is connected through a gateway to the network in the cloud. The address of the user's gateway must be preconfigured on the VPN server in the cloud. Note that although L2TP-over-IPsec can be used to set up Site-to-Site VPNs, this is not the primary intent of this feature. For more information, see [Seção 15.17.4, "Configurando uma conexão VPN Site-to-Site"](#)

15.17.1. Configurando VPN

Para configurar VPN para a nuvem:

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, clique em Global Settings.
3. Configure os seguintes parâmetros de configuração global.
 - `remote.access.vpn.client.ip.range` – The range of IP addresses to be allocated to remote access VPN clients. The first IP in the range is used by the VPN server.
 - `remote.access.vpn.psk.length` – Tamanho da chave IPsec.
 - `remote.access.vpn.user.limit` – Número máximo de usuários VPN por conta.

Para habilitar VPN para uma rede em particular:

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, clique em Network.
3. Clique no nome da rede com a qual você quer trabalhar.
4. Clique em View IP Addresses.
5. Clique em um dos endereços IP apresentados.
6.  Click the Enable VPN button.

A chave IPsec é apresentada em uma janela popup.

15.17.2. Usando VPN com Windows

O procedimento para usar VPN varia de acordo com a versão do Windows. Geralmente, o usuário deve editar as propriedades da VPN e certificar-se de que a rota default não é a VPN. Os passos seguintes são para clientes Windows L2TP no Windows Vista. Os comandos devem ser similares para outras versões do Windows.

1. Faça login na interface de usuário do CloudStack e clique no IP NAT de origem para a conta. A aba VPN deve exibir a chave IPsec pré-compartilhada. Tome nota disto e do IP NAT de origem. A interface de usuário também lista um ou mais usuários e suas senhas. Escolha um destes usuários ou, se nenhum existe, adicione um usuário e uma senha.
2. No Windows, vá para o Painel de Controle, selecione Centro de Rede e Compartilhamento. Clique em Configurar uma conexão ou uma rede.
3. No próximo diálogo, selecione Conectar a um local de trabalho.
4. No próximo diálogo, selecione Use minha conexão com a Internet (VPN).
5. In the next dialog, enter the source NAT IP from step [1](#) and give the connection a name. Check Don't connect now.
6. In the next dialog, enter the user name and password selected in step [1](#).
7. Clique em Create.

8. Volte ao painel de Controle e clique em Conexões de Rede para ver a nova conexão. A conexão ainda não está ativa.
9. Clique com o botão da direita na nova conexão e selecione Propriedades. No diálogo de Propriedades, selecione a aba Rede.
10. In Type of VPN, choose L2TP IPsec VPN, then click IPsec settings. Select Use preshared key. Enter the preshared key from Step 1.
11. A conexão está pronta para ser ativada. Volte a Painel de Controle -> Conexões de Rede e dê dois cliques na conexão criada.
12. Enter the user name and password from Step 1.

15.17.3. Using VPN with Mac OS X

First, be sure you've configured the VPN settings in your CloudStack install. This section is only concerned with connecting via Mac OS X to your VPN.

Note, these instructions were written on Mac OS X 10.7.5. They may differ slightly in older or newer releases of Mac OS X.

1. On your Mac, open System Preferences and click Network.
2. Make sure Send all traffic over VPN connection is not checked.
3. If your preferences are locked, you'll need to click the lock in the bottom left-hand corner to make any changes and provide your administrator credentials.
4. You will need to create a new network entry. Click the plus icon on the bottom left-hand side and you'll see a dialog that says "Select the interface and enter a name for the new service." Select VPN from the Interface drop-down menu, and "L2TP over IPsec" for the VPN Type. Enter whatever you like within the "Service Name" field.
5. You'll now have a new network interface with the name of whatever you put in the "Service Name" field. For the purposes of this example, we'll assume you've named it "CloudStack." Click on that interface and provide the IP address of the interface for your VPN under the Server Address field, and the user name for your VPN under Account Name.
6. Click Authentication Settings, and add the user's password under User Authentication and enter the pre-shared IPsec key in the Shared Secret field under Machine Authentication. Click OK.
7. You may also want to click the "Show VPN status in menu bar" but that's entirely optional.
8. Now click "Connect" and you will be connected to the CloudStack VPN.

15.17.4. Configurando uma conexão VPN Site-to-Site

Uma conexão VPN Site-to-Site o ajuda a estabelecer uma conexão segura de um centro de dados empresarial à infraestrutura de nuvem. Isto permite o acesso de usuários a máquinas virtuais hóspedes estabelecendo uma conexão VPN de um equipamento no centro de dados da empresa ao roteador virtual da conta. Possuindo este recurso é eliminada a necessidade de estabelecer conexões VPN a máquinas virtuais individuais.

Os terminais suportados nos centros de dados remotos são:

- Cisco ISR com IOS 12.4 ou posterior

- Roteadores Juniper J-Series com JunOS 9.5 ou posterior



Nota

Adicionalmente aos equipamentos Cisco e Juniper específicos listados acima, a expectativa é que qualquer equipamento Cisco ou Juniper executando os sistemas operacionais suportados sejam capazes de estabelecer conexões VPN.

Para configurar uma conexão VPN Site-to-Site, execute o seguinte:

1. Crie uma Virtual Private Cloud (VPC).
Veja [Seção 15.19, “Configuring a Virtual Private Cloud”](#).
2. Crie um VPN Customer Gateway.
3. Crie um VPN gateway para a VPC que você criou.
4. Crie uma conexão VPN do VPN gateway da VPC para o customer VPN gateway.



Nota

Appropriate events are generated on the CloudStack UI when status of a Site-to-Site VPN connection changes from connected to disconnected, or vice versa. Currently no events are generated when establishing a VPN connection fails or pending.

15.17.4.1. Creating and Updating a VPN Customer Gateway



Nota

A VPN customer gateway can be connected to only one VPN gateway at a time.

To add a VPN Customer Gateway:

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. In the Select view, select VPN Customer Gateway.
4. Click Add site-to-site VPN.

add VPN Customer Gateway

* Name:

* Gateway:

* CIDR list:

* IPsec Preshared-Key:

IKE Encryption:

IKE Hash:

IKE DH:

ESP Encryption:

ESP Hash:

Perfect Forward Secrecy:

IKE lifetime (second):

ESP Lifetime (second):

Dead Peer Detection:

Cancel OK

Forneça as seguintes informações:

- **Name:** A unique name for the VPN customer gateway you create.
- **Gateway:** The IP address for the remote gateway.
- **CIDR list:** The guest CIDR list of the remote subnets. Enter a CIDR or a comma-separated list of CIDRs. Ensure that a guest CIDR list is not overlapped with the VPC's CIDR, or another guest CIDR. The CIDR must be RFC1918-compliant.
- **IPsec Preshared Key:** Preshared keying is a method where the endpoints of the VPN share a secret key. This key value is used to authenticate the customer gateway and the VPC VPN gateway to each other.



Nota

The IKE peers (VPN end points) authenticate each other by computing and sending a keyed hash of data that includes the Preshared key. If the receiving peer is able to create the same hash independently by using its Preshared key, it knows that both peers must share the same secret, thus authenticating the customer gateway.

- **IKE Encryption:** The Internet Key Exchange (IKE) policy for phase-1. The supported encryption algorithms are AES128, AES192, AES256, and 3DES. Authentication is accomplished through the Preshared Keys.



Nota

The phase-1 is the first phase in the IKE process. In this initial negotiation phase, the two VPN endpoints agree on the methods to be used to provide security for the underlying IP traffic. The phase-1 authenticates the two VPN gateways to each other, by confirming that the remote gateway has a matching Preshared Key.

- **IKE Hash:** The IKE hash for phase-1. The supported hash algorithms are SHA1 and MD5.
- **IKE DH:** A public-key cryptography protocol which allows two parties to establish a shared secret over an insecure communications channel. The 1536-bit Diffie-Hellman group is used within IKE to establish session keys. The supported options are None, Group-5 (1536-bit) and Group-2 (1024-bit).
- **ESP Encryption:** Encapsulating Security Payload (ESP) algorithm within phase-2. The supported encryption algorithms are AES128, AES192, AES256, and 3DES.



Nota

The phase-2 is the second phase in the IKE process. The purpose of IKE phase-2 is to negotiate IPsec security associations (SA) to set up the IPsec tunnel. In phase-2, new keying material is extracted from the Diffie-Hellman key exchange in phase-1, to provide session keys to use in protecting the VPN data flow.

- **ESP Hash:** Encapsulating Security Payload (ESP) hash for phase-2. Supported hash algorithms are SHA1 and MD5.
- **Perfect Forward Secrecy:** Perfect Forward Secrecy (or PFS) is the property that ensures that a session key derived from a set of long-term public and private keys will not be compromised. This property enforces a new Diffie-Hellman key exchange. It provides the keying material

that has greater key material life and thereby greater resistance to cryptographic attacks. The available options are None, Group-5 (1536-bit) and Group-2 (1024-bit). The security of the key exchanges increase as the DH groups grow larger, as does the time of the exchanges.



Nota


When PFS is turned on, for every negotiation of a new phase-2 SA the two gateways must generate a new set of phase-1 keys. This adds an extra layer of protection that PFS adds, which ensures if the phase-2 SA's have expired, the keys used for new phase-2 SA's have not been generated from the current phase-1 keying material.

- **IKE Lifetime (seconds):** The phase-1 lifetime of the security association in seconds. Default is 86400 seconds (1 day). Whenever the time expires, a new phase-1 exchange is performed.
 - **ESP Lifetime (seconds):** The phase-2 lifetime of the security association in seconds. Default is 3600 seconds (1 hour). Whenever the value is exceeded, a re-key is initiated to provide a new IPsec encryption and authentication session keys.
 - **Dead Peer Detection:** A method to detect an unavailable Internet Key Exchange (IKE) peer. Select this option if you want the virtual router to query the liveliness of its IKE peer at regular intervals. It's recommended to have the same configuration of DPD on both side of VPN connection.
5. Clique em OK.

Updating and Removing a VPN Customer Gateway

You can update a customer gateway either with no VPN connection, or related VPN connection is in error state.

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. In the Select view, select VPN Customer Gateway.
4. Select the VPN customer gateway you want to work with.

5. To modify the required parameters, click the Edit VPN Customer Gateway button 

6. To remove the VPN customer gateway, click the Delete VPN Customer Gateway button 

7. Clique em OK.

15.17.4.2. Criando um gateway VPN para o VPC

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Na view Select, selecione VPC.

Todas as VPCs que você criou para a conta são listadas na página.

4. Clique no botão Configure da VPC na qual você deseja implementar as máquinas virtuais.

Na página VPC apresentada, todas as camadas que você criou são listadas em um diagrama.

5. Clique no ícone Settings.

As seguintes opções são apresentadas.

- Endereços IP
- Gateways
- Site-to-Site VPN
- Network ACLs

6. Selecione Site-to-Site VPN.

Se você está criando o gateway VPN pela primeira vez, selecionando Site-to-Site VPN requer de você criar um gateway VPN.

7. No diálogo de confirmação, clique Yes para confirmar.

Em poucos momentos, o gateway VPN é criado. Você será solicitado ver os detalhes do gateway VPN que você criou. Clique Yes para confirmar.

Os seguintes detalhes são apresentados na página VPN Gateway:

- Endereço IP
- Conta
- Domínio

15.17.4.3. Criando uma conexão VPN

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Na view Select, selecione VPC.

Todas as VPCs que você cria para a conta são listadas na página.

4. Clique no botão Configure da VPC na qual você deseja implementar as máquinas virtuais.

Na página VPC apresentada, todas as camadas que você criou são listadas em um diagrama.

5. Clique no ícone Settings.

As seguintes opções são apresentadas.

- IP Addresses
- Gateways
- Site-to-Site VPN

- Network ACLs


6. Selecione Site-to-Site VPN.

A página Site-to-Site VPN é apresentada.

7. Na lista suspensa Select View, certifique-se de que VPN Connection é selecionada.

8. Clique em Create VPN Connection.

O diálogo Create VPN Connection é apresentado:



9. Selecione o gateway desejado, então clique OK para confirmar.

Em poucos momentos, a conexão VPN é apresentada.

A seguinte informação sobre a conexão VPN é apresentada:

- Endereço IP
- Gateway
- Estado
- Chave IPsec pré compartilhada
- Política IKE
- Política ESP

15.17.4.4. Reiniciando e removendo uma conexão VPN

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.

2. Na barra de navegação à esquerda, selecione Network.

3. Na view Select, selecione VPC.

Todas as VPCs que você criou para a conta são listadas na página.

4. Clique no botão Configure da VPC na qual você deseja implementar as máquinas virtuais.

Na página VPC apresentada, todas as camadas que você criou são listadas em um diagrama.

5. Clique no ícone Settings.

As seguintes opções são apresentadas.

- Endereços IP
- Gateways
- Site-to-Site VPN
- Network ACLs

6. Selecione Site-to-Site VPN.


A página Site-to-Site VPN é apresentada.

7. Na lista suspensa Select View, certifique-se de que VPN Connection é selecionada.

Todas as conexões VPN que você criou são apresentadas.

8. Selecione a conexão VPN com a qual você quer trabalhar.

A aba Details é apresentada.

9. Para remover uma conexão VPN, clique no botão Delete VPN connection 

Para reiniciar uma conexão VPN, clique no botão Reset VPN connection presente na aba Details.



15.18. About Inter-VLAN Routing

Inter-VLAN Routing is the capability to route network traffic between VLANs. This feature enables you to build Virtual Private Clouds (VPC), an isolated segment of your cloud, that can hold multi-tier applications. These tiers are deployed on different VLANs that can communicate with each other. You provision VLANs to the tiers you create, and VMs can be deployed on different tiers. The VLANs are connected to a virtual router, which facilitates communication between the VMs. In effect, you can segment VMs by means of VLANs into different networks that can host multi-tier applications, such as Web, Application, or Database. Such segmentation by means of VLANs logically separate application VMs for higher security and lower broadcasts, while remaining physically connected to the same device.

This feature is supported on XenServer and VMware hypervisors.

The major advantages are:

- The administrator can deploy a set of VLANs and allow users to deploy VMs on these VLANs. A guest VLAN is randomly allotted to an account from a pre-specified set of guest VLANs. All the VMs of a certain tier of an account reside on the guest VLAN allotted to that account.



Nota

A VLAN allocated for an account cannot be shared between multiple accounts.

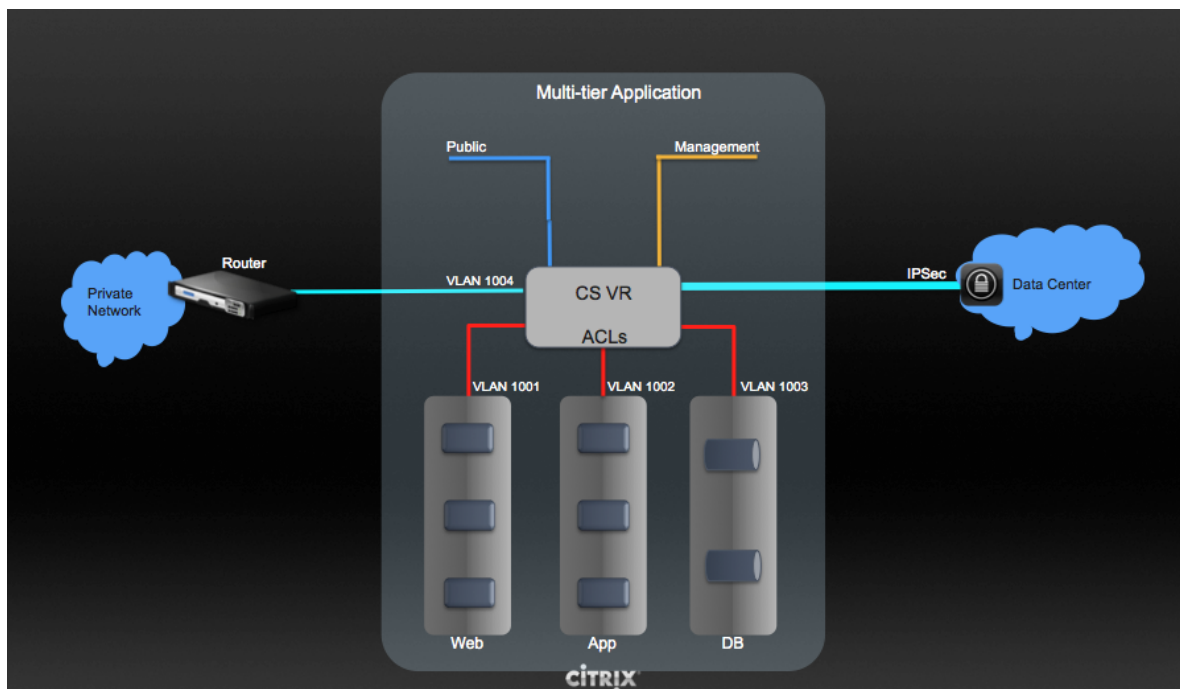
- The administrator can allow users create their own VPC and deploy the application. In this scenario, the VMs that belong to the account are deployed on the VLANs allotted to that account.

- Both administrators and users can create multiple VPCs. The guest network NIC is plugged to the VPC virtual router when the first VM is deployed in a tier.
- The administrator can create the following gateways to send to or receive traffic from the VMs:
 - **VPN Gateway:** For more information, see [Seção 15.17.4.2, “Criando um gateway VPN para o VPC”](#).
 - **Public Gateway:** The public gateway for a VPC is added to the virtual router when the virtual router is created for VPC. The public gateway is not exposed to the end users. You are not allowed to list it, nor allowed to create any static routes.
 - **Private Gateway:** For more information, see [Seção 15.19.5, “Adicionando um gateway privado a uma VPC”](#).
- Both administrators and users can create various possible destinations-gateway combinations. However, only one gateway of each type can be used in a deployment.

For example:

- **VLANs and Public Gateway:** For example, an application is deployed in the cloud, and the Web application VMs communicate with the Internet.
- **VLANs, VPN Gateway, and Public Gateway:** For example, an application is deployed in the cloud; the Web application VMs communicate with the Internet; and the database VMs communicate with the on-premise devices.
- The administrator can define Access Control List (ACL) on the virtual router to filter the traffic among the VLANs or between the Internet and a VLAN. You can define ACL based on CIDR, port range, protocol, type code (if ICMP protocol is selected) and Ingress/Egress type.

The following figure shows the possible deployment scenarios of a Inter-VLAN setup:



To set up a multi-tier Inter-VLAN deployment, see [Seção 15.19, “Configuring a Virtual Private Cloud”](#).

15.19. Configuring a Virtual Private Cloud

15.19.1. About Virtual Private Clouds

CloudStack Virtual Private Cloud is a private, isolated part of CloudStack. A VPC can have its own virtual network topology that resembles a traditional physical network. You can launch VMs in the virtual network that can have private addresses in the range of your choice, for example: 10.0.0.0/16. You can define network tiers within your VPC network range, which in turn enables you to group similar kinds of instances based on IP address range.

For example, if a VPC has the private range 10.0.0.0/16, its guest networks can have the network ranges 10.0.1.0/24, 10.0.2.0/24, 10.0.3.0/24, and so on.

Major Components of a VPC:

A VPC is comprised of the following network components:

- **VPC:** A VPC acts as a container for multiple isolated networks that can communicate with each other via its virtual router.
- **Network Tiers:** Each tier acts as an isolated network with its own VLANs and CIDR list, where you can place groups of resources, such as VMs. The tiers are segmented by means of VLANs. The NIC of each tier acts as its gateway.
- **Virtual Router:** A virtual router is automatically created and started when you create a VPC. The virtual router connect the tiers and direct traffic among the public gateway, the VPN gateways, and the NAT instances. For each tier, a corresponding NIC and IP exist in the virtual router. The virtual router provides DNS and DHCP services through its IP.
- **Public Gateway:** The traffic to and from the Internet routed to the VPC through the public gateway. In a VPC, the public gateway is not exposed to the end user; therefore, static routes are not support for the public gateway.
- **Private Gateway:** All the traffic to and from a private network routed to the VPC through the private gateway. For more information, see [Seção 15.19.5, “Adicionando um gateway privado a uma VPC”](#).
- **VPN Gateway:** The VPC side of a VPN connection.
- **Site-to-Site VPN Connection:** A hardware-based VPN connection between your VPC and your datacenter, home network, or co-location facility. For more information, see [Seção 15.17.4, “Configurando uma conexão VPN Site-to-Site”](#).
- **Customer Gateway:** The customer side of a VPN Connection. For more information, see [Seção 15.17.4.1, “Creating and Updating a VPN Customer Gateway”](#).
- **NAT Instance:** An instance that provides Port Address Translation for instances to access the Internet via the public gateway. For more information, see [Seção 15.19.9, “Habilitando ou desabilitando NAT estática em uma VPC”](#).

Network Architecture in a VPC

In a VPC, the following four basic options of network architectures are present:

- VPC with a public gateway only
- VPC with public and private gateways

- VPC with public and private gateways and site-to-site VPN access
- VPC with a private gateway only and site-to-site VPN access

Connectivity Options for a VPC

You can connect your VPC to:

- The Internet through the public gateway.
- The corporate datacenter by using a site-to-site VPN connection through the VPN gateway.
- Both the Internet and your corporate datacenter by using both the public gateway and a VPN gateway.

VPC Network Considerations

Consider the following before you create a VPC:

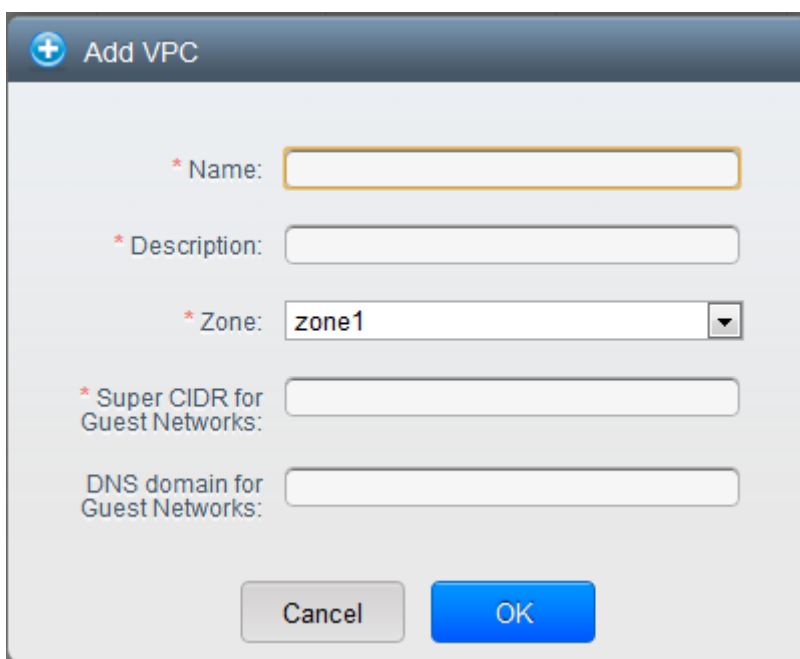
- A VPC, by default, is created in the enabled state.
- A VPC can be created in Advance zone only, and can't belong to more than one zone at a time.
- The default number of VPCs an account can create is 20. However, you can change it by using the `max.account.vpcs` global parameter, which controls the maximum number of VPCs an account is allowed to create.
- The default number of tiers an account can create within a VPC is 3. You can configure this number by using the `vpc.max.networks` parameter.
- Each tier should have a unique CIDR in the VPC. Ensure that the tier's CIDR should be within the VPC CIDR range.
- A tier belongs to only one VPC.
- All network tiers inside the VPC should belong to the same account.
- When a VPC is created, by default, a SourceNAT IP is allocated to it. The Source NAT IP is released only when the VPC is removed.
- A public IP can be used for only one purpose at a time. If the IP is a sourceNAT, it cannot be used for StaticNAT or port forwarding.
- The instances only have a private IP address that you provision. To communicate with the Internet, enable NAT to an instance that you launch in your VPC.
- Only new networks can be added to a VPC. The maximum number of networks per VPC is limited by the value you specify in the `vpc.max.networks` parameter. The default value is three.
- The load balancing service can be supported by only one tier inside the VPC.
- If an IP address is assigned to a tier:
 - That IP can't be used by more than one tier at a time in the VPC. For example, if you have tiers A and B, and a public IP1, you can create a port forwarding rule by using the IP either for A or B, but not for both.
 - That IP can't be used for StaticNAT, load balancing, or port forwarding rules for another guest network inside the VPC.

- Remote access VPN is not supported in VPC networks.

15.19.2. Adding a Virtual Private Cloud

When creating the VPC, you simply provide the zone and a set of IP addresses for the VPC network address space. You specify this set of addresses in the form of a Classless Inter-Domain Routing (CIDR) block.

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Na view Select, selecione VPC.
4. Click Add VPC. The Add VPC page is displayed as follows:



Forneça as seguintes informações:

- **Name:** A short name for the VPC that you are creating.
- **Description:** A brief description of the VPC.
- **Zone:** Choose the zone where you want the VPC to be available.
- **Super CIDR for Guest Networks:** Defines the CIDR range for all the tiers (guest networks) within a VPC. When you create a tier, ensure that its CIDR is within the Super CIDR value you enter. The CIDR must be RFC1918 compliant.
- **DNS domain for Guest Networks:** If you want to assign a special domain name, specify the DNS suffix. This parameter is applied to all the tiers within the VPC. That implies, all the tiers you create in the VPC belong to the same DNS domain. If the parameter is not specified, a DNS domain name is generated automatically.

15.19.3. Adding Tiers

Tiers are distinct locations within a VPC that act as isolated networks, which do not have access to other tiers by default. Tiers are set up on different VLANs that can communicate with each other by

using a virtual router. Tiers provide inexpensive, low latency network connectivity to other tiers within the VPC.

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Na view Select, selecione VPC.

All the VPC that you have created for the account is listed in the page.



Nota

The end users can see their own VPCs, while root and domain admin can see any VPC they are authorized to see.

4. Click the Configure button of the VPC for which you want to set up tiers.

The Add new tier dialog is displayed, as follows:

The screenshot shows a dialog box titled "Add new tier". It contains the following fields and values:

- * Name: NewTier
- * Network Offering: DefaultIsolatedNetworkOfferingForV (dropdown menu)
- * Gateway: 10.0.0.1
- * Netmask: 255.255.255.0

At the bottom of the dialog, there are two buttons: "Cancel" and "OK".

If you have already created tiers, the VPC diagram is displayed. Click Create Tier to add a new tier.

5. Especifique o seguinte:

All the fields are mandatory.

- **Name:** A unique name for the tier you create.
- **Network Offering:** The following default network offerings are listed: DefaultIsolatedNetworkOfferingForVpcNetworksNoLB, DefaultIsolatedNetworkOfferingForVpcNetworks

In a VPC, only one tier can be created by using LB-enabled network offering.

- **Gateway:** The gateway for the tier you create. Ensure that the gateway is within the Super CIDR range that you specified while creating the VPC, and is not overlapped with the CIDR of any existing tier within the VPC.
- **Netmask:** The netmask for the tier you create.

For example, if the VPC CIDR is 10.0.0.0/16 and the network tier CIDR is 10.0.1.0/24, the gateway of the tier is 10.0.1.1, and the netmask of the tier is 255.255.255.0.

6. Clique em OK.
7. Continue with configuring access control list for the tier.

15.19.4. Configuring Access Control List

Define Network Access Control List (ACL) on the VPC virtual router to control incoming (ingress) and outgoing (egress) traffic between the VPC tiers, and the tiers and Internet. By default, all incoming and outgoing traffic to the guest networks is blocked. To open the ports, you must create a new network ACL. The network ACLs can be created for the tiers only if the NetworkACL service is supported.

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Na view Select, selecione VPC.

Todas as VPCs que você criou para a conta são listadas na página.

4. Clique no ícone Settings.

As seguintes opções são apresentadas.

- Endereços IP
- Gateways
- Site-to-Site VPN
- Network ACLs

5. Select Network ACLs.


The Network ACLs page is displayed.

6. Click Add Network ACLs.

To add an ACL rule, fill in the following fields to specify what kind of network traffic is allowed in this tier.

- **CIDR:** The CIDR acts as the Source CIDR for the Ingress rules, and Destination CIDR for the Egress rules. To accept traffic only from or to the IP addresses within a particular address block, enter a CIDR or a comma-separated list of CIDRs. The CIDR is the base IP address of the incoming traffic. For example, 192.168.0.0/22. To allow all CIDRs, set to 0.0.0.0/0.
- **Protocol:** The networking protocol that sources use to send traffic to the tier. The TCP and UDP protocols are typically used for data exchange and end-user communications. The ICMP protocol is typically used to send error messages or network monitoring data.

- **Start Port, End Port** (TCP, UDP only): A range of listening ports that are the destination for the incoming traffic. If you are opening a single port, use the same number in both fields.
- **Select Tier**: Select the tier for which you want to add this ACL rule.
- **ICMP Type, ICMP Code** (ICMP only): The type of message and error code that will be sent.
- **Traffic Type**: Select the traffic type you want to apply.
 - **Egress**: To add an egress rule, select Egress from the Traffic type drop-down box and click Add. This specifies what type of traffic is allowed to be sent out of VM instances in this tier. If no egress rules are specified, all traffic from the tier is allowed out at the VPC virtual router. Once egress rules are specified, only the traffic specified in egress rules and the responses to any traffic that has been allowed in through an ingress rule are allowed out. No egress rule is required for the VMs in a tier to communicate with each other.
 - **Ingress**: To add an ingress rule, select Ingress from the Traffic type drop-down box and click Add. This specifies what network traffic is allowed into the VM instances in this tier. If no ingress rules are specified, then no traffic will be allowed in, except for responses to any traffic that has been allowed out through an egress rule.

 **Nota**

By default, all incoming and outgoing traffic to the guest networks is blocked. To open the ports, create a new network ACL.

7. Click Add. The ACL rule is added.

To view the list of ACL rules you have added, click the desired tier from the Network ACLs page, then select the Network ACL tab.

Network Details		Network ACL		IP Addresses				
CIDR	Protocol	Start Port	End Port	ICMP Type	ICMP Code	Traffic type	Add rule	Actions
<input type="text"/>	TCP	<input type="text"/>	<input type="text"/>			Ingress	<input type="button" value="Add"/>	
0.0.0.0/0	TCP	1	65535			Ingress		
0.0.0.0/0	TCP	1	65535			Egress		
0.0.0.0/0	ICMP			-1	-1	Egress		
0.0.0.0/0	ICMP			-1	-1	Ingress		

You can edit the tags assigned to the ACL rules and delete the ACL rules you have created. Click the appropriate button in the Actions column.

15.19.5. Adicionando um gateway privado a uma VPC

Um gateway privado somente pode ser adicionado pelo administrador root. A rede privada da VPC tem uma relação 1:1 com a NIC da rede física. Nenhum gateway com VLAN e IP duplicado é permitido no mesmo centro de dados.

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Na view Select, selecione VPC.

Todas as VPCs que você criou para a conta são listadas na página.

4. Clique no botão Configure da VPC na qual você deseja configurar regras de balanceamento de carga.

Na página VPC apresentada, todas as camadas que você criou são listadas em um diagrama.

5. Clique no ícone Settings.

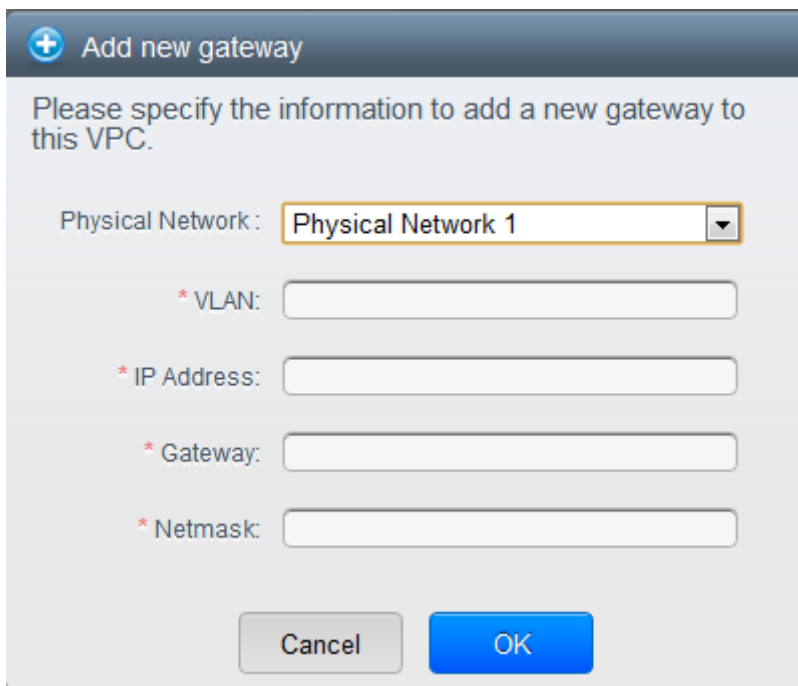
As seguintes opções são apresentadas.

- Endereços IP
- Private Gateways
- Site-to-Site VPN
- Network ACLs

6. Selecione Private Gateways.

A página Gateways é apresentada.

7. Clique em Add new gateway:



8. Especifique o seguinte:

- **Physical Network:** A rede física que você criou na zona.
- **IP Address:** O endereço IP associado com o gateway VPC.
- **Gateway:** O gateway através do qual o tráfego é roteado de e para a VPC.
- **Netmask:** A máscara de rede associada com o gateway VPC.
- **VLAN:** A VLAN associada com o gateway VPC.

O novo gateway aparece na lista. Você pode repetir estes passos para adicionar mais gateways para esta VPC.

15.19.6. Implantando máquinas virtuais na camada

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Na view Select, selecione VPC.

Todas as VPCs que você criou para a conta são listadas na página.

4. Clique no botão Configure da VPC na qual você deseja implementar as máquinas virtuais.

Na página VPC apresentada, todas as camadas que você criou são listadas.

5. Clique no botão Add VM da camada a qual você deseja adicionar uma máquina virtual.

A página Add Instance é apresentada.

Siga as instruções na tela para adicionar uma instância. Para informações sobre adição de uma instância, veja a seção Adicionando instâncias no Guia de Instalação.

15.19.7. Obtendo um novo endereço IP para uma VPC

Quando você obtém um endereço IP, todos os endereços IP são alocados à VPC, não às redes hóspedes na VPC. Os IPs são associados à rede hóspede somente quando a primeira regra de port-forwarding, balanceamento de carga ou NAT estática é criada para o IP ou para a rede. Um IP não pode ser associado a mais de uma rede de cada vez.

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Na view Select, selecione VPC.

Todas as VPCs que você criou para a conta são listadas na página.

4. Clique no botão Configure da VPC na qual você deseja implementar as máquinas virtuais.

Na página VPC apresentada, todas as camadas que você criou são listadas em um diagrama.

5. Clique no ícone Settings.

As seguintes opções são apresentadas.

- Endereços IP
- Gateways
- Site-to-Site VPN
- Network ACLs

6. Selecione IP Addresses.

A página IP Addresses é apresentada.

7. Clique em Acquire New IP, e clique Yes no diálogo de confirmação.

Você deve confirmar porque, tipicamente, endereços IP são um recurso limitado. Em alguns momentos, o novo endereço IP deve aparecer no estado Allocated. Agora você pode usar o endereço IP no encaminhamento de porta, balanceamento de carga, ou regras NAT estáticas.

15.19.8. Liberando um endereço IP atribuído a uma VPC

O endereço IP é um recurso limitado. Se você não precisa mais de um IP particular, você pode desassociá-lo de sua VPC e retorná-lo ao pool de endereços disponíveis. Um endereço IP pode ser liberado de sua camada somente quando todas as rede de rede (port forwarding, balanceamento de carga, ou NAT estática) são removidas para este endereço IP. O endereço IP liberado ainda irá pertencer à mesma VPC.

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Na view Select, selecione VPC.

Todas as VPCs que você criou para a conta são listadas na página.

4. Clique no botão Configure da VPC da qual você deseja liberar o IP.

Na página VPC apresentada, todas as camadas que você criou são listadas em um diagrama.

5. Clique no ícone Settings.


As seguintes opções são apresentadas.

- Endereços IP
- Gateways
- Site-to-Site VPN
- Network ACLs

6. Selecione IP Addresses.

A página IP Addresses é apresentada.

7. Clique no IP que você deseja liberar.

8. Na aba Details, clique no botão Release IP 

15.19.9. Habilitando ou desabilitando NAT estática em uma VPC

Uma regra NAT estática mapeia um endereço IP público para o endereço IP privado de uma máquina virtual em uma VPC para permitir tráfego da Internet para ela. Esta seção informa como habilitar ou desabilitar NAT estática para um endereço IP em particular em uma VPC.

Se regras de encaminhamento de portas já estão em efeito para um endereço IP, você não pode habilitar NAT estática para este IP.

Se uma máquina virtual hóspede faz parte de mais de uma rede, regras de NAT estática funcionarão somente se elas estão definidas na rede default.

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Na view Select, selecione VPC.

Todas as VPCs que você criou para a conta são listadas na página.

4. Clique no botão Configure da VPC na qual você deseja implementar as máquinas virtuais.

Na página VPC apresentada, todas as camadas que você criou são listadas em um diagrama.

5. Clique no ícone Settings.


As seguintes opções são apresentadas.

- Endereços IP
- Gateways
- Site-to-Site VPN
- Network ACLs

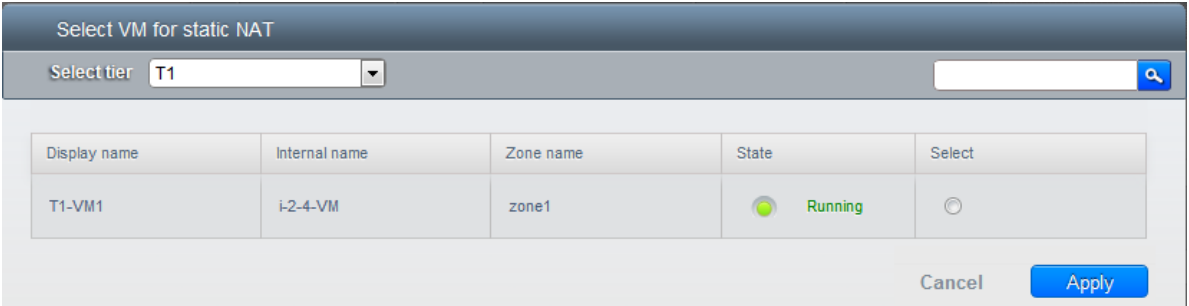
6. Selecione IP Addresses.

A página IP Addresses é apresentada.

7. Clique no IP com o qual você deseja trabalhar.

8. Na aba Details, clique no botão Static NAT.  O botão alterna entre Enable e Disable, dependendo se NAT estática está habilitada ou não para o endereço IP.

9. Se você está habilitando NAT estática, um diálogo é apresentado como se segue:



Display name	Internal name	Zone name	State	Select
T1-VM1	i-2-4-VM	zone1	● Running	<input type="radio"/>

10. Selecione a camada e a máquina virtual de destino, e então clique em Apply.

15.19.10. Adicionando regras de balanceamento de carga em uma VPC

Um usuário ou administrador do CloudStack pode criar regras de balanceamento de carga que distribuem o tráfego recebido em um endereço IP público por uma ou mais máquinas virtuais que pertencem a uma camada de rede que provê serviço de balanceamento de carga em uma VPC. Um usuário cria uma regra, especifica um algoritmo e assinala a regra a um conjunto de máquinas virtuais em uma VPC.

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Na view Select, selecione VPC.

Todas as VPCs que você criou para a conta são listadas na página.

4. Clique no botão Configure da VPC na qual você deseja configurar regras de balanceamento de carga.

Na página VPC apresentada, todas as camadas que você criou são listadas em um diagrama.

5. Clique no ícone Settings.

As seguintes opções são apresentadas.

- IP Addresses
- Gateways
- Site-to-Site VPN
- Network ACLs

6. Selecione IP Addresses.

A página IP Addresses é apresentada.

7. Clique no endereço IP para o qual você deseja criar a regra, então clique na aba Configuration.
8. No nó Load Balancing do diagrama, clique em View All.
9. Selecione a camada para a qual você deseja aplicar a regra.



Nota

Em uma VPC, o serviço de balanceamento de carga é suportado somente em uma única camada.

10. Especifique o seguinte:

- **Name:** Um nome para a regra de balanceamento de carga.
- **Public Port:** A porta que recebe o tráfego entrante a ser balanceado.

- **Private Port:** A porta que as máquinas virtuais usarão para receber o tráfego.
- **Algorithm.** Escolha o algoritmo de balanceamento de carga que você deseja que o CloudStack use. O CloudStack suporta os seguintes algoritmos bem conhecidos:
 - Round-robin
 - Menos conexões
 - Origem
- **Stickiness.** (Opcional) Clique Configure e escolha o algoritmo para a política de afinidade. Veja Sticky Session Policies for Load Balancer Rules.
- **Add VMs:** Clique em Add VMs, então selecione duas ou mais máquinas virtuais que irão dividir a carga do tráfego entrante, e clique em Apply.

A nova regra de balanceamento de carga aparece na lista. Você pode repetir estes passos para adicionar mais regras de balanceamento de carga para este endereço IP.

15.19.11. Adicionando uma regra de encaminhamento de porta em uma VPC

1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Na view Select, selecione VPC.

Todas as VPCs que você criou para a conta são listadas na página.

4. Clique no botão Configure da VPC na qual você deseja implementar as máquinas virtuais.

Na página VPC apresentada, todas as camadas que você criou são listadas em um diagrama.

5. Clique no ícone Settings.

As seguintes opções são apresentadas.

- Endereços IP
- Gateways
- Site-to-Site VPN
- Network ACLs

6. Escolha um endereço IP existente ou obtenha um novo endereço IP. Clique no nome do endereço IP na lista.

A página IP Addresses é apresentada.

7. Clique no endereço IP para o qual você deseja criar a regra, então clique na aba Configuration.
8. No nó Port Forwarding do diagrama, clique em View All.
9. Selecione a camada para a qual você deseja aplicar a regra.

10. Especifique o seguinte:

- **Public Port:** A porta para a qual tráfego público será encaminhado no endereço IP que você obteve no passo anterior.
- **Private Port:** A porta na qual a instância está ouvindo por tráfego público encaminhado.
- **Protocol:** O protocolo de comunicações em uso entre as duas portas.
 - TCP
 - UDP
- **Add VM:** Clique em Add VM. Selecione o nome da instância a qual esta regra se aplica, e clique Apply.

Você pode testar a regra abrindo uma sessão ssh com a instância.

15.19.12. Removing Tiers

You can remove a tier from a VPC. A removed tier cannot be revoked. When a tier is removed, only the resources of the tier are expunged. All the network rules (port forwarding, load balancing and staticNAT) and the IP addresses associated to the tier are removed. The IP address still be belonging to the same VPC.


1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Na view Select, selecione VPC.

All the VPC that you have created for the account is listed in the page.

4. Click the Configure button of the VPC for which you want to set up tiers.

The Configure VPC page is displayed. Locate the tier you want to work with.

5. Click the Remove VPC button:



Wait for some time for the tier to be removed.

15.19.13. Editing, Restarting, and Removing a Virtual Private Cloud




Nota


Ensure that all the tiers are removed before you remove a VPC.


1. Faça login na interface de usuário do CloudStack como administrador ou usuário final.
2. Na barra de navegação à esquerda, selecione Network.
3. Na view Select, selecione VPC.

Todas as VPCs que você criou para a conta são listadas na página.

4. Select the VPC you want to work with.

5. To remove, click the Remove VPC button 

You can edit the name and description of a VPC. To do that, select the VPC, then click the Edit button. 

To restart a VPC, select the VPC, then click the Restart button. 

15.20. Persistent Networks

The network that you can provision without having to deploy any VMs on it is called a persistent network. A persistent network can be part of a VPC or a non-VPC environment.

When you create other types of network, a network is only a database entry until the first VM is created on that network. When the first VM is created, a VLAN ID is assigned and the network is provisioned. Also, when the last VM is destroyed, the VLAN ID is released and the network is no longer available. With the addition of persistent network, you will have the ability to create a network in CloudStack in which physical devices can be deployed without having to run any VMs. Additionally, you can deploy physical devices on that network.

One of the advantages of having a persistent network is that you can create a VPC with a tier consisting of only physical devices. For example, you might create a VPC for a three-tier application, deploy VMs for Web and Application tier, and use physical machines for the Database tier. Another use case is that if you are providing services by using physical hardware, you can define the network as persistent and therefore even if all its VMs are destroyed the services will not be discontinued.

15.20.1. Persistent Network Considerations

- Persistent network is designed for isolated networks.
- All default network offerings are non-persistent.
- A network offering cannot be editable because changing it affects the behavior of the existing networks that were created using this network offering.

- When you create a guest network, the network offering that you select defines the network persistence. This in turn depends on whether persistent network is enabled in the selected network offering.
- An existing network can be made persistent by changing its network offering to an offering that has the Persistent option enabled. While setting this property, even if the network has no running VMs, the network is provisioned.
- An existing network can be made non-persistent by changing its network offering to an offering that has the Persistent option disabled. If the network has no running VMs, during the next network garbage collection run the network is shut down.
- When the last VM on a network is destroyed, the network garbage collector checks if the network offering associated with the network is persistent, and shuts down the network only if it is non-persistent.

15.20.2. Creating a Persistent Guest Network

To create a persistent network, perform the following:

1. Create a network offering with the Persistent option enabled.
See [Seção 9.4.1, “Creating a New Network Offering”](#).
2. Select Network from the left navigation pane.
3. Select the guest network that you want to offer this network service to.
4. Click the Edit button.
5. From the Network Offering drop-down, select the persistent network offering you have just created.
6. Clique em OK.

Working with System Virtual Machines

CloudStack uses several types of system virtual machines to perform tasks in the cloud. In general CloudStack manages these system VMs and creates, starts, and stops them as needed based on scale and immediate needs. However, the administrator should be aware of them and their roles to assist in debugging issues.



Nota

You can configure the `system.vm.random.password` parameter to create a random system VM password to ensure higher security. If you reset the value for `system.vm.random.password` to true and restart the Management Server, a random password is generated and stored encrypted in the database. You can view the decrypted password under the `system.vm.password` global parameter on the CloudStack UI or by calling the `listConfigurations` API.

16.1. The System VM Template

The System VMs come from a single template. The System VM has the following characteristics:

- Debian 6.0 ("Squeeze"), 2.6.32 kernel with the latest security patches from the Debian security APT repository
- Has a minimal set of packages installed thereby reducing the attack surface
- 32-bit for enhanced performance on Xen/VMWare
- pvops kernel with Xen PV drivers, KVM virtio drivers, and VMware tools for optimum performance on all hypervisors
- Xen tools inclusion allows performance monitoring
- Latest versions of HAProxy, iptables, IPsec, and Apache from debian repository ensures improved security and speed
- Latest version of JRE from Sun/Oracle ensures improved security and speed

16.2. Multiple System VM Support for VMware

Every CloudStack zone has single System VM for template processing tasks such as downloading templates, uploading templates, and uploading ISOs. In a zone where VMware is being used, additional System VMs can be launched to process VMware-specific tasks such as taking snapshots and creating private templates. The CloudStack management server launches additional System VMs for VMware-specific tasks as the load increases. The management server monitors and weights all commands sent to these System VMs and performs dynamic load balancing and scaling-up of more System VMs.

16.3. Console Proxy

The Console Proxy is a type of System Virtual Machine that has a role in presenting a console view via the web UI. It connects the user's browser to the VNC port made available via the hypervisor for the console of the guest. Both the administrator and end user web UIs offer a console connection.

Clicking a console icon brings up a new window. The AJAX code downloaded into that window refers to the public IP address of a console proxy VM. There is exactly one public IP address allocated per console proxy VM. The AJAX application connects to this IP. The console proxy then proxies the connection to the VNC port for the requested VM on the Host hosting the guest.



Nota

The hypervisors will have many ports assigned to VNC usage so that multiple VNC sessions can occur simultaneously.

There is never any traffic to the guest virtual IP, and there is no need to enable VNC within the guest.

The console proxy VM will periodically report its active session count to the Management Server. The default reporting interval is five seconds. This can be changed through standard Management Server configuration with the parameter `consoleproxy.loadscan.interval`.

Assignment of guest VM to console proxy is determined by first determining if the guest VM has a previous session associated with a console proxy. If it does, the Management Server will assign the guest VM to the target Console Proxy VM regardless of the load on the proxy VM. Failing that, the first available running Console Proxy VM that has the capacity to handle new sessions is used.

Console proxies can be restarted by administrators but this will interrupt existing console sessions for users.

16.3.1. Using a SSL Certificate for the Console Proxy

The console viewing functionality uses a dynamic DNS service under the domain name `realhostip.com` to assist in providing SSL security to console sessions. The console proxy is assigned a public IP address. In order to avoid browser warnings for mismatched SSL certificates, the URL for the new console window is set to the form of `https://aaa-bbb-ccc-ddd.realhostip.com`. You will see this URL during console session creation. CloudStack includes the `realhostip.com` SSL certificate in the console proxy VM. Of course, CloudStack cannot know about the DNS A records for our customers' public IPs prior to shipping the software. CloudStack therefore runs a dynamic DNS server that is authoritative for the `realhostip.com` domain. It maps the `aaa-bbb-ccc-ddd` part of the DNS name to the IP address `aaa.bbb.ccc.ddd` on lookups. This allows the browser to correctly connect to the console proxy's public IP, where it then expects and receives a SSL certificate for `realhostip.com`, and SSL is set up without browser warnings.

16.3.2. Changing the Console Proxy SSL Certificate and Domain

If the administrator prefers, it is possible for the URL of the customer's console session to show a domain other than `realhostip.com`. The administrator can customize the displayed domain by selecting a different domain and uploading a new SSL certificate and private key. The domain must run a DNS service that is capable of resolving queries for addresses of the form `aaa-bbb-ccc-ddd.your.domain` to an IPv4 IP address in the form `aaa.bbb.ccc.ddd`, for example, `202.8.44.1`. To change the console proxy domain, SSL certificate, and private key:

1. Set up dynamic name resolution or populate all possible DNS names in your public IP range into your existing DNS server with the format `aaa-bbb-ccc-ddd.company.com -> aaa.bbb.ccc.ddd`.
2. Generate the private key and certificate signing request (CSR). When you are using `openssl` to generate private/public key pairs and CSRs, for the private key that you are going to paste into the CloudStack UI, be sure to convert it into PKCS#8 format.

- a. Generate a new 2048-bit private key

```
openssl genrsa -des3 -out yourprivate.key 2048
```

- b. Generate a new certificate CSR

```
openssl req -new -key yourprivate.key -out yourcertificate.csr
```

- c. Head to the website of your favorite trusted Certificate Authority, purchase an SSL certificate, and submit the CSR. You should receive a valid certificate in return

- d. Convert your private key format into PKCS#8 encrypted format.

```
openssl pkcs8 -topk8 -in yourprivate.key -out yourprivate.pkcs8.encrypted.key
```

- e. Convert your PKCS#8 encrypted private key into the PKCS#8 format that is compliant with CloudStack

```
openssl pkcs8 -in yourprivate.pkcs8.encrypted.key -out yourprivate.pkcs8.key
```

3. In the Update SSL Certificate screen of the CloudStack UI, paste the following:

- The certificate you've just generated.
- The private key you've just generated.
- The desired new domain name; for example, company.com

4. The desired new domain name; for example, company.com

This stops all currently running console proxy VMs, then restarts them with the new certificate and key. Users might notice a brief interruption in console availability.

The Management Server generates URLs of the form "aaa-bbb-ccc-ddd.company.com" after this change is made. The new console requests will be served with the new DNS domain name, certificate, and key.

16.4. Virtual Router

The virtual router is a type of System Virtual Machine. The virtual router is one of the most frequently used service providers in CloudStack. The end user has no direct access to the virtual router. Users can ping the virtual router and take actions that affect it (such as setting up port forwarding), but users do not have SSH access into the virtual router.

There is no mechanism for the administrator to log in to the virtual router. Virtual routers can be restarted by administrators, but this will interrupt public network access and other services for end users. A basic test in debugging networking issues is to attempt to ping the virtual router from a guest VM. Some of the characteristics of the virtual router are determined by its associated system service offering..

16.4.1. Configuring the Virtual Router

You can set the following:

- IP range
- Supported network services
- Default domain name for the network serviced by the virtual router
- Gateway IP address
- How often CloudStack fetches network usage statistics from CloudStack virtual routers. If you want to collect traffic metering data from the virtual router, set the global configuration parameter `router.stats.interval`. If you are not using the virtual router to gather network usage statistics, set it to 0.

16.4.2. Upgrading a Virtual Router with System Service Offerings

When CloudStack creates a virtual router, it uses default settings which are defined in a default system service offering. See [Seção 8.2, “System Service Offerings”](#). All the virtual routers in a single guest network use the same system service offering. You can upgrade the capabilities of the virtual router by creating and applying a custom system service offering.

1. Define your custom system service offering. See [Seção 8.2.1, “Creating a New System Service Offering”](#). In System VM Type, choose Domain Router.
2. Associate the system service offering with a network offering. See [Seção 9.4.1, “Creating a New Network Offering”](#).
3. Apply the network offering to the network where you want the virtual routers to use the new system service offering. If this is a new network, follow the steps in Adding an Additional Guest Network on page 66. To change the service offering for existing virtual routers, follow the steps in [Seção 15.6.2, “Alterando a oferta de rede em uma rede hóspede”](#).

16.4.3. Best Practices for Virtual Routers

- **WARNING:** Restarting a virtual router from a hypervisor console deletes all the iptables rules. To work around this issue, stop the virtual router and start it from the CloudStack UI.
- **WARNING:** Do not use the `destroyRouter` API when only one router is available in the network, because `restartNetwork` API with the `cleanup=false` parameter can't recreate it later. If you want to destroy and recreate the single router available in the network, use the `restartNetwork` API with the `cleanup=true` parameter.

16.5. Secondary Storage VM

In addition to the hosts, CloudStack's Secondary Storage VM mounts and writes to secondary storage.

Submissions to secondary storage go through the Secondary Storage VM. The Secondary Storage VM can retrieve templates and ISO images from URLs using a variety of protocols.

The secondary storage VM provides a background task that takes care of a variety of secondary storage activities: downloading a new template to a Zone, copying templates between Zones, and snapshot backups.

The administrator can log in to the secondary storage VM if needed.

System Reliability and High Availability

17.1. HA for Management Server

The CloudStack Management Server should be deployed in a multi-node configuration such that it is not susceptible to individual server failures. The Management Server itself (as distinct from the MySQL database) is stateless and may be placed behind a load balancer.

Normal operation of Hosts is not impacted by an outage of all Management Servers. All guest VMs will continue to work.

When the Management Server is down, no new VMs can be created, and the end user and admin UI, API, dynamic load distribution, and HA will cease to work.

17.2. Management Server Load Balancing

CloudStack can use a load balancer to provide a virtual IP for multiple Management Servers. The administrator is responsible for creating the load balancer rules for the Management Servers. The application requires persistence or stickiness across multiple sessions. The following chart lists the ports that should be load balanced and whether or not persistence is required.

Even if persistence is not required, enabling it is permitted.

Source Port	Destination Port	Protocolo	Persistence Required?
80 or 443	8080 (or 20400 with AJP)	HTTP (or AJP)	Sim
8250	8250	TCP	Sim
8096	8096	HTTP	Não

In addition to above settings, the administrator is responsible for setting the 'host' global config value from the management server IP to load balancer virtual IP address. If the 'host' value is not set to the VIP for Port 8250 and one of your management servers crashes, the UI is still available but the system VMs will not be able to contact the management server.

17.3. HA-Enabled Virtual Machines

The user can specify a virtual machine as HA-enabled. By default, all virtual router VMs and Elastic Load Balancing VMs are automatically configured as HA-enabled. When an HA-enabled VM crashes, CloudStack detects the crash and restarts the VM automatically within the same Availability Zone. HA is never performed across different Availability Zones. CloudStack has a conservative policy towards restarting VMs and ensures that there will never be two instances of the same VM running at the same time. The Management Server attempts to start the VM on another Host in the same cluster.

HA features work with iSCSI or NFS primary storage. HA with local storage is not supported.

17.4. HA for Hosts

The user can specify a virtual machine as HA-enabled. By default, all virtual router VMs and Elastic Load Balancing VMs are automatically configured as HA-enabled. When an HA-enabled VM crashes,

CloudStack detects the crash and restarts the VM automatically within the same Availability Zone. HA is never performed across different Availability Zones. CloudStack has a conservative policy towards restarting VMs and ensures that there will never be two instances of the same VM running at the same time. The Management Server attempts to start the VM on another Host in the same cluster.

HA features work with iSCSI or NFS primary storage. HA with local storage is not supported.

17.4.1. Dedicated HA Hosts

One or more hosts can be designated for use only by HA-enabled VMs that are restarting due to a host failure. Setting up a pool of such dedicated HA hosts as the recovery destination for all HA-enabled VMs is useful to:

- Make it easier to determine which VMs have been restarted as part of the CloudStack high-availability function. If a VM is running on a dedicated HA host, then it must be an HA-enabled VM whose original host failed. (With one exception: It is possible for an administrator to manually migrate any VM to a dedicated HA host.)
- Keep HA-enabled VMs from restarting on hosts which may be reserved for other purposes.

The dedicated HA option is set through a special host tag when the host is created. To allow the administrator to dedicate hosts to only HA-enabled VMs, set the global configuration variable `ha.tag` to the desired tag (for example, "ha_host"), and restart the Management Server. Enter the value in the Host Tags field when adding the host(s) that you want to dedicate to HA-enabled VMs.



Nota

If you set `ha.tag`, be sure to actually use that tag on at least one host in your cloud. If the tag specified in `ha.tag` is not set for any host in the cloud, the HA-enabled VMs will fail to restart after a crash.

17.5. Primary Storage Outage and Data Loss

When a primary storage outage occurs the hypervisor immediately stops all VMs stored on that storage device. Guests that are marked for HA will be restarted as soon as practical when the primary storage comes back on line. With NFS, the hypervisor may allow the virtual machines to continue running depending on the nature of the issue. For example, an NFS hang will cause the guest VMs to be suspended until storage connectivity is restored. Primary storage is not designed to be backed up. Individual volumes in primary storage can be backed up using snapshots.

17.6. Secondary Storage Outage and Data Loss

For a Zone that has only one secondary storage server, a secondary storage outage will have feature level impact to the system but will not impact running guest VMs. It may become impossible to create a VM with the selected template for a user. A user may also not be able to save snapshots or examine/restore saved snapshots. These features will automatically be available when the secondary storage comes back online.

Secondary storage data loss will impact recently added user data including templates, snapshots, and ISO images. Secondary storage should be backed up periodically. Multiple secondary storage servers can be provisioned within each zone to increase the scalability of the system.

17.7. Limiting the Rate of API Requests

You can limit the rate at which API requests can be placed for each account. This is useful to avoid malicious attacks on the Management Server, prevent performance degradation, and provide fairness to all accounts.

If the number of API calls exceeds the threshold, an error message is returned for any additional API calls. The caller will have to retry these API calls at another time.

17.7.1. Configuring the API Request Rate

To control the API request rate, use the following global configuration settings:

- `api.throttling.enabled` - Enable/Disable API throttling. By default, this setting is false, so API throttling is not enabled.
- `api.throttling.interval` (in seconds) - Time interval during which the number of API requests is to be counted. When the interval has passed, the API count is reset to 0.
- `api.throttling.max` - Maximum number of APIs that can be placed within the `api.throttling.interval` period.
- `api.throttling.cachesize` - Cache size for storing API counters. Use a value higher than the total number of accounts managed by the cloud. One cache entry is needed for each account, to store the running API total for that account.

17.7.2. Limitations on API Throttling

The following limitations exist in the current implementation of this feature.



Nota

Even with these limitations, CloudStack is still able to effectively use API throttling to avoid malicious attacks causing denial of service.

- In a deployment with multiple Management Servers, the cache is not synchronized across them. In this case, CloudStack might not be able to ensure that only the exact desired number of API requests are allowed. In the worst case, the number of API calls that might be allowed is (number of Management Servers) * (`api.throttling.max`).
- The API commands `resetApiLimit` and `getApiLimit` are limited to the Management Server where the API is invoked.

Gerenciando a nuvem

18.1. Using Tags to Organize Resources in the Cloud

A tag is a key-value pair that stores metadata about a resource in the cloud. Tags are useful for categorizing resources. For example, you can tag a user VM with a value that indicates the user's city of residence. In this case, the key would be "city" and the value might be "Toronto" or "Tokyo." You can then request CloudStack to find all resources that have a given tag; for example, VMs for users in a given city.

You can tag a user virtual machine, volume, snapshot, guest network, template, ISO, firewall rule, port forwarding rule, public IP address, security group, load balancer rule, project, VPC, network ACL, or static route. You can not tag a remote access VPN.

You can work with tags through the UI or through the API commands `createTags`, `deleteTags`, and `listTags`. You can define multiple tags for each resource. There is no limit on the number of tags you can define. Each tag can be up to 255 characters long. Users can define tags on the resources they own, and administrators can define tags on any resources in the cloud.

An optional input parameter, "tags," exists on many of the list* API commands. The following example shows how to use this new parameter to find all the volumes having tag `region=canada` OR tag `city=Toronto`:

```
command=listVolumes
  &listAll=true
  &tags[0].key=region
  &tags[0].value=canada
  &tags[1].key=city
  &tags[1].value=Toronto
```

The following API commands have the "tags" input parameter:

- `listVirtualMachines`
- `listVolumes`
- `listSnapshots`
- `listNetworks`
- `listTemplates`
- `listIsos`
- `listFirewallRules`
- `listPortForwardingRules`
- `listPublicIpAddresses`
- `listSecurityGroups`
- `listLoadBalancerRules`
- `listProjects`
- `listVPCs`

- listNetworkACLs
- listStaticRoutes

18.2. Changing the Database Configuration

The CloudStack Management Server stores database configuration information (e.g., hostname, port, credentials) in the file `/etc/cloud/management/db.properties`. To effect a change, edit this file on each Management Server, then restart the Management Server.

18.3. Changing the Database Password

You may need to change the password for the MySQL account used by CloudStack. If so, you'll need to change the password in MySQL, and then add the encrypted password to `/etc/cloud/management/db.properties`.

1. Before changing the password, you'll need to stop CloudStack's management server and the usage engine if you've deployed that component.

```
# service cloudstack-management stop
# service cloudstack-usage stop
```

2. Next, you'll update the password for the CloudStack user on the MySQL server.

```
# mysql -u root -p
```

At the MySQL shell, you'll change the password and flush privileges:

```
update mysql.user set password=PASSWORD("newpassword123") where User='cloud';
flush privileges;
quit;
```

3. The next step is to encrypt the password and copy the encrypted password to CloudStack's database configuration (`/etc/cloud/management/db.properties`).

```
# java -classpath /usr/share/java/cloud-jasypt-1.8.jar \
org.jasypt.intf.cli.JasyptPBEStrEncryptionCLI encrypt.sh \ input="newpassword123"
password="`cat /etc/cloud/management/key`" \ verbose=false
```



File encryption type

Note that this is for the file encryption type. If you're using the web encryption type then you'll use `password="management_server_secret_key"`

4. Now, you'll update `/etc/cloud/management/db.properties` with the new ciphertext. Open `/etc/cloud/management/db.properties` in a text editor, and update these parameters:

```
db.cloud.password=ENC(encrypted_password_from_above)
db.usage.password=ENC(encrypted_password_from_above)
```

5. After copying the new password over, you can now start CloudStack (and the usage engine, if necessary).

```
# service cloudstack-management start
# service cloudstack-usage start
```

18.4. Administrator Alerts

The system provides alerts and events to help with the management of the cloud. Alerts are notices to an administrator, generally delivered by e-mail, notifying the administrator that an error has occurred in the cloud. Alert behavior is configurable.

Events track all of the user and administrator actions in the cloud. For example, every guest VM start creates an associated event. Events are stored in the Management Server's database.

Emails will be sent to administrators under the following circumstances:

- The Management Server cluster runs low on CPU, memory, or storage resources
- The Management Server loses heartbeat from a Host for more than 3 minutes
- The Host cluster runs low on CPU, memory, or storage resources

18.5. Customizing the Network Domain Name

The root administrator can optionally assign a custom DNS suffix at the level of a network, account, domain, zone, or entire CloudStack installation, and a domain administrator can do so within their own domain. To specify a custom domain name and put it into effect, follow these steps.

1. Set the DNS suffix at the desired scope
 - At the network level, the DNS suffix can be assigned through the UI when creating a new network, as described in [Seção 15.6.1, “Adicionando uma rede hóspede adicional”](#) or with the `updateNetwork` command in the CloudStack API.
 - At the account, domain, or zone level, the DNS suffix can be assigned with the appropriate CloudStack API commands: `createAccount`, `editAccount`, `createDomain`, `editDomain`, `createZone`, or `editZone`.
 - At the global level, use the configuration parameter `guest.domain.suffix`. You can also use the CloudStack API command `updateConfiguration`. After modifying this global configuration, restart the Management Server to put the new setting into effect.
2. To make the new DNS suffix take effect for an existing network, call the CloudStack API command `updateNetwork`. This step is not necessary when the DNS suffix was specified while creating a new network.

The source of the network domain that is used depends on the following rules.

- For all networks, if a network domain is specified as part of a network's own configuration, that value is used.

- For an account-specific network, the network domain specified for the account is used. If none is specified, the system looks for a value in the domain, zone, and global configuration, in that order.
- For a domain-specific network, the network domain specified for the domain is used. If none is specified, the system looks for a value in the zone and global configuration, in that order.
- For a zone-specific network, the network domain specified for the zone is used. If none is specified, the system looks for a value in the global configuration.

18.6. Stopping and Restarting the Management Server

The root administrator will need to stop and restart the Management Server from time to time.

For example, after changing a global configuration parameter, a restart is required. If you have multiple Management Server nodes, restart all of them to put the new parameter value into effect consistently throughout the cloud..

To stop the Management Server, issue the following command at the operating system prompt on the Management Server node:

```
# service cloudstack-management stop
```

To start the Management Server:

```
# service cloudstack-management start
```

To stop the Management Server:

```
# service cloudstack-management stop
```

Global Configuration Parameters

19.1. Setting Global Configuration Parameters

CloudStack provides parameters that you can set to control many aspects of the cloud. When CloudStack is first installed, and periodically thereafter, you might need to modify these settings.

1. Log in to the UI as administrator.
2. Na barra de navegação à esquerda, clique em Global Settings.
3. In Select View, choose one of the following:
 - Global Settings. This displays a list of the parameters with brief descriptions and current values.
 - Hypervisor Capabilities. This displays a list of hypervisor versions with the maximum number of guests supported for each.
4. Use the search box to narrow down the list to those you are interested in.
5. Click the Edit icon to modify a value. If you are viewing Hypervisor Capabilities, you must click the name of the hypervisor first to display the editing screen.

19.2. About Global Configuration Parameters

CloudStack provides a variety of settings you can use to set limits, configure features, and enable or disable features in the cloud. Once your Management Server is running, you might need to set some of these global configuration parameters, depending on what optional features you are setting up.

To modify global configuration parameters, use the steps in "Setting Global Configuration Parameters."

The documentation for each CloudStack feature should direct you to the names of the applicable parameters. Many of them are discussed in the CloudStack Administration Guide. The following table shows a few of the more useful parameters.

Field	Valor
management.network.cidr	A CIDR that describes the network that the management CIDRs reside on. This variable must be set for deployments that use vSphere. It is recommended to be set for other deployments as well. Example: 192.168.3.0/24.
xen.setup.multipath	For XenServer nodes, this is a true/false variable that instructs CloudStack to enable iSCSI multipath on the XenServer Hosts when they are added. This defaults to false. Set it to true if you would like CloudStack to enable multipath.

Field	Valor
	<p>If this is true for a NFS-based deployment multipath will still be enabled on the XenServer host. However, this does not impact NFS operation and is harmless.</p>
<p>secstorage.allowed.internal.sites</p>	<p>This is used to protect your internal network from rogue attempts to download arbitrary files using the template download feature. This is a comma-separated list of CIDRs. If a requested URL matches any of these CIDRs the Secondary Storage VM will use the private network interface to fetch the URL. Other URLs will go through the public interface. We suggest you set this to 1 or 2 hardened internal machines where you keep your templates. For example, set it to 192.168.1.66/32.</p>
<p>use.local.storage</p>	<p>Determines whether CloudStack will use storage that is local to the Host for data disks, templates, and snapshots. By default CloudStack will not use this storage. You should change this to true if you want to use local storage and you understand the reliability and feature drawbacks to choosing local storage.</p>
<p>host</p>	<p>This is the IP address of the Management Server. If you are using multiple Management Servers you should enter a load balanced IP address that is reachable via the private network.</p>
<p>default.page.size</p>	<p>Maximum number of items per page that can be returned by a CloudStack API command. The limit applies at the cloud level and can vary from cloud to cloud. You can override this with a lower value on a particular API call by using</p>

Field	Valor
	the page and pagesize API command parameters. For more information, see the Developer's Guide. Default: 500.
ha.tag	The label you want to use throughout the cloud to designate certain hosts as dedicated HA hosts. These hosts will be used only for HA-enabled VMs that are restarting due to the failure of another host. For example, you could set this to ha_host. Specify the ha.tag value as a host tag when you add a new host to the cloud.

CloudStack API

The CloudStack API is a low level API that has been used to implement the CloudStack web UIs. It is also a good basis for implementing other popular APIs such as EC2/S3 and emerging DMTF standards.

Many CloudStack API calls are asynchronous. These will return a Job ID immediately when called. This Job ID can be used to query the status of the job later. Also, status calls on impacted resources will provide some indication of their state.

The API has a REST-like query basis and returns results in XML or JSON.

See [the Developer's Guide](#)¹ and [the API Reference](#)².

20.1. Provisioning and Authentication API

CloudStack expects that a customer will have their own user provisioning infrastructure. It provides APIs to integrate with these existing systems where the systems call out to CloudStack to add/remove users..

CloudStack supports pluggable authenticators. By default, CloudStack assumes it is provisioned with the user's password, and as a result authentication is done locally. However, external authentication is possible as well. For example, see [Using an LDAP Server for User Authentication](#).

20.2. Allocators

CloudStack enables administrators to write custom allocators that will choose the Host to place a new guest and the storage host from which to allocate guest virtual disk images.

20.3. User Data and Meta Data

CloudStack provides API access to attach user data to a deployed VM. Deployed VMs also have access to instance metadata via the virtual router.

User data can be accessed once the IP address of the virtual router is known. Once the IP address is known, use the following steps to access the user data:

1. Run the following command to find the virtual router.

```
# cat /var/lib/dhclient/dhclient-eth0.leases | grep dhcp-server-identifier | tail -1
```

2. Access user data by running the following command using the result of the above command

```
# curl http://10.1.1.1/latest/user-data
```

Meta Data can be accessed similarly, using a URL of the form `http://10.1.1.1/latest/meta-data/{metadata type}`. (For backwards compatibility, the previous URL `http://10.1.1.1/latest/{metadata type}` is also supported.) For metadata type, use one of the following:

- `service-offering`. A description of the VMs service offering

¹ http://docs.cloudstack.org/CloudStack_Documentation/Developer's_Guide%3A_CloudStack

² http://docs.cloudstack.org/CloudStack_Documentation/API_Reference%3A_CloudStack

- availability-zone. The Zone name
- local-ipv4. The guest IP of the VM
- local-hostname. The hostname of the VM
- public-ipv4. The first public IP for the router. (E.g. the first IP of eth2)
- public-hostname. This is the same as public-ipv4
- instance-id. The instance name of the VM

Tuning

This section provides tips on how to improve the performance of your cloud.

21.1. Performance Monitoring

Host and guest performance monitoring is available to end users and administrators. This allows the user to monitor their utilization of resources and determine when it is appropriate to choose a more powerful service offering or larger disk.

21.2. Increase Management Server Maximum Memory

If the Management Server is subject to high demand, the default maximum JVM memory allocation can be insufficient. To increase the memory:

1. Edit the Tomcat configuration file:

```
/etc/cloud/management/tomcat6.conf
```

2. Change the command-line parameter `-XmxNNNm` to a higher value of `N`.

For example, if the current value is `-Xmx128m`, change it to `-Xmx1024m` or higher.

3. To put the new setting into effect, restart the Management Server.

```
# service cloudstack-management restart
```

For more information about memory issues, see "FAQ: Memory" at [Tomcat Wiki](#).¹

21.3. Set Database Buffer Pool Size

It is important to provide enough memory space for the MySQL database to cache data and indexes:

1. Edit the MySQL configuration file:

```
/etc/my.cnf
```

2. Insert the following line in the `[mysqld]` section, below the `datadir` line. Use a value that is appropriate for your situation. We recommend setting the buffer pool at 40% of RAM if MySQL is on the same server as the management server or 70% of RAM if MySQL has a dedicated server. The following example assumes a dedicated server with 1024M of RAM.

```
innodb_buffer_pool_size=700M
```

3. Reinicie o serviço MySQL.

```
# service mysqld restart
```

¹ <http://wiki.apache.org/tomcat/FAQ/Memory>

For more information about the buffer pool, see "The InnoDB Buffer Pool" at [MySQL Reference Manual](#)².

21.4. Set and Monitor Total VM Limits per Host

The CloudStack administrator should monitor the total number of VM instances in each cluster, and disable allocation to the cluster if the total is approaching the maximum that the hypervisor can handle. Be sure to leave a safety margin to allow for the possibility of one or more hosts failing, which would increase the VM load on the other hosts as the VMs are automatically redeployed. Consult the documentation for your chosen hypervisor to find the maximum permitted number of VMs per host, then use CloudStack global configuration settings to set this as the default limit. Monitor the VM activity in each cluster at all times. Keep the total number of VMs below a safe level that allows for the occasional host failure. For example, if there are N hosts in the cluster, and you want to allow for one host in the cluster to be down at any given time, the total number of VM instances you can permit in the cluster is at most $(N-1) * (\text{per-host-limit})$. Once a cluster reaches this number of VMs, use the CloudStack UI to disable allocation of more VMs to the cluster.

21.5. Configure XenServer dom0 Memory

Configure the XenServer dom0 settings to allocate more memory to dom0. This can enable XenServer to handle larger numbers of virtual machines. We recommend 2940 MB of RAM for XenServer dom0. For instructions on how to do this, see [Citrix Knowledgebase Article](#)³. The article refers to XenServer 5.6, but the same information applies to XenServer 6

² <http://dev.mysql.com/doc/refman/5.5/en/innodb-buffer-pool.html>

³ <http://support.citrix.com/article/CTX126531>

Troubleshooting

22.1. Events

An event is essentially a significant or meaningful change in the state of both virtual and physical resources associated with a cloud environment. Events are used by monitoring systems, usage and billing systems, or any other event-driven workflow systems to discern a pattern and make the right business decision. In CloudStack an event could be a state change of virtual or physical resources, an action performed by an user (action events), or policy based events (alerts).

22.1.1. Event Logs

There are two types of events logged in the CloudStack Event Log. Standard events log the success or failure of an event and can be used to identify jobs or processes that have failed. There are also long running job events. Events for asynchronous jobs log when a job is scheduled, when it starts, and when it completes. Other long running synchronous jobs log when a job starts, and when it completes. Long running synchronous and asynchronous event logs can be used to gain more information on the status of a pending job or can be used to identify a job that is hanging or has not started. The following sections provide more information on these events..

22.1.2. Event Notification

Event notification framework provides a means for the Management Server components to publish and subscribe to CloudStack events. Event notification is achieved by implementing the concept of event bus abstraction in the Management Server. An event bus is introduced in the Management Server that allows the CloudStack components and extension plug-ins to subscribe to the events by using the Advanced Message Queuing Protocol (AMQP) client. In CloudStack, a default implementation of event bus is provided as a plug-in that uses the RabbitMQ AMQP client. The AMQP client pushes the published events to a compatible AMQP server. Therefore all the CloudStack events are published to an exchange in the AMQP server.

A new event for state change, resource state change, is introduced as part of Event notification framework. Every resource, such as user VM, volume, NIC, network, public IP, snapshot, and template, is associated with a state machine and generates events as part of the state change. That implies that a change in the state of a resource results in a state change event, and the event is published in the corresponding state machine on the event bus. All the CloudStack events (alerts, action events, usage events) and the additional category of resource state change events, are published on to the events bus.

Use Cases

The following are some of the use cases:

- Usage or Billing Engines: A third-party cloud usage solution can implement a plug-in that can connect to CloudStack to subscribe to CloudStack events and generate usage data. The usage data is consumed by their usage software.
- AMQP plug-in can place all the events on the a message queue, then a AMQP message broker can provide topic-based notification to the subscribers.
- Publish and Subscribe notification service can be implemented as a pluggable service in CloudStack that can provide rich set of APIs for event notification, such as topics-based subscription and notification. Additionally, the pluggable service can deal with multi-tenancy, authentication, and authorization issues.

Configuração

As a CloudStack administrator, perform the following one-time configuration to enable event notification framework. At run time no changes can control the behaviour.

1. Open '**componentContext.xml**'.
2. Define a bean named **eventNotificationBus** as follows:
 - name : Specify a name for the bean.
 - server : The name or the IP address of the RabbitMQ AMQP server.
 - port : The port on which RabbitMQ server is running.
 - username : The username associated with the account to access the RabbitMQ server.
 - password : The password associated with the username of the account to access the RabbitMQ server.
 - exchange : The exchange name on the RabbitMQ server where CloudStack events are published.

A sample bean is given below:

```
<bean id="eventNotificationBus"
  class="org.apache.cloudstack.mom.rabbitmq.RabbitMQEventBus">
  <property name="name" value="eventNotificationBus"/>
  <property name="server" value="127.0.0.1"/>
  <property name="port" value="5672"/>
  <property name="username" value="guest"/>
  <property name="password" value="guest"/>
  <property name="exchange" value="cloudstack-events"/>
</bean>
```

The **eventNotificationBus** bean represents the **org.apache.cloudstack.mom.rabbitmq.RabbitMQEventBus** class.

3. Reinicie o servidor de gerenciamento.

22.1.3. Standard Events

The events log records three types of standard events.

- INFO. This event is generated when an operation has been successfully performed.
- WARN. This event is generated in the following circumstances.
 - When a network is disconnected while monitoring a template download.
 - When a template download is abandoned.
 - When an issue on the storage server causes the volumes to fail over to the mirror storage server.
- ERROR. This event is generated when an operation has not been successfully performed

22.1.4. Long Running Job Events

The events log records three types of standard events.

- INFO. This event is generated when an operation has been successfully performed.
- WARN. This event is generated in the following circumstances.
 - When a network is disconnected while monitoring a template download.
 - When a template download is abandoned.
 - When an issue on the storage server causes the volumes to fail over to the mirror storage server.
- ERROR. This event is generated when an operation has not been successfully performed

22.1.5. Event Log Queries

Database logs can be queried from the user interface. The list of events captured by the system includes:

- Virtual machine creation, deletion, and on-going management operations
- Virtual router creation, deletion, and on-going management operations
- Template creation and deletion
- Network/load balancer rules creation and deletion
- Storage volume creation and deletion
- User login and logout

22.2. Working with Server Logs

The CloudStack Management Server logs all web site, middle tier, and database activities for diagnostics purposes in `/var/log/cloudstack/management/`. The CloudStack logs a variety of error messages. We recommend this command to find the problematic output in the Management Server log:



Nota

Ao copiar e colar um comando, certifique-se que o comando tenha colado como uma única linha antes de executar. Alguns viewers documento pode introduzir quebras de linha indesejadas no texto copiado.

```
grep -i -E 'exception|unable|fail|invalid|leak|warn|error' /var/log/cloudstack/management/management-server.log
```

The CloudStack processes requests with a Job ID. If you find an error in the logs and you are interested in debugging the issue you can grep for this job ID in the management server log. For example, suppose that you find the following ERROR message:

```
2010-10-04 13:49:32,595 ERROR [cloud.vm.UserVmManagerImpl] (Job-Executor-11:job-1076) Unable to find any host for [User|i-8-42-VM-untagged]
```

Note that the job ID is 1076. You can track back the events relating to job 1076 with the following grep:

```
grep "job-1076)" management-server.log
```

The CloudStack Agent Server logs its activities in `/var/log/cloudstack/agent/`.

22.3. Data Loss on Exported Primary Storage

Symptom

Loss of existing data on primary storage which has been exposed as a Linux NFS server export on an iSCSI volume.

Cause

It is possible that a client from outside the intended pool has mounted the storage. When this occurs, the LVM is wiped and all data in the volume is lost

Solution

When setting up LUN exports, restrict the range of IP addresses that are allowed access by specifying a subnet mask. For example:

```
echo "/export 192.168.1.0/24(rw,async,no_root_squash)" > /etc/exports
```

Adjust the above command to suit your deployment needs.

More Information

See the export procedure in the "Secondary Storage" section of the CloudStack Installation Guide

22.4. Recovering a Lost Virtual Router

Symptom

A virtual router is running, but the host is disconnected. A virtual router no longer functions as expected.

Cause

The Virtual router is lost or down.

Solution

If you are sure that a virtual router is down forever, or no longer functions as expected, destroy it. You must create one afresh while keeping the backup router up and running (it is assumed this is in a redundant router setup):

- Force stop the router. Use the stopRouter API with forced=true parameter to do so.
- Before you continue with destroying this router, ensure that the backup router is running. Otherwise the network connection will be lost.

- Destroy the router by using the destroyRouter API.

Recreate the missing router by using the restartNetwork API with cleanup=false parameter. For more information about redundant router setup, see [Creating a New Network Offering](#).

For more information about the API syntax, see the API Reference at http://docs.cloudstack.org/CloudStack_Documentation/API_Reference%3A_CloudStack API Reference.

22.5. Maintenance mode not working on vCenter

Symptom

Host was placed in maintenance mode, but still appears live in vCenter.

Cause

The CloudStack administrator UI was used to place the host in scheduled maintenance mode. This mode is separate from vCenter's maintenance mode.

Solution

Use vCenter to place the host in maintenance mode.

More Information

See [Seção 11.2, "Scheduled Maintenance and Maintenance Mode for Hosts"](#)

22.6. Unable to deploy VMs from uploaded vSphere template

Symptom

When attempting to create a VM, the VM will not deploy.

Cause

If the template was created by uploading an OVA file that was created using vSphere Client, it is possible the OVA contained an ISO image. If it does, the deployment of VMs from the template will fail.

Solution

Remove the ISO and re-upload the template.

22.7. Unable to power on virtual machine on VMware

Symptom

Virtual machine does not power on. You might see errors like:

- Unable to open Swap File

- Unable to access a file since it is locked
- Unable to access Virtual machine configuration

Cause

A known issue on VMware machines. ESX hosts lock certain critical virtual machine files and file systems to prevent concurrent changes. Sometimes the files are not unlocked when the virtual machine is powered off. When a virtual machine attempts to power on, it can not access these critical files, and the virtual machine is unable to power on.

Solution

See the following:

[VMware Knowledge Base Article](#)¹

22.8. Load balancer rules fail after changing network offering

Symptom

After changing the network offering on a network, load balancer rules stop working.

Cause

Load balancing rules were created while using a network service offering that includes an external load balancer device such as NetScaler, and later the network service offering changed to one that uses the CloudStack virtual router.

Solution

Create a firewall rule on the virtual router for each of your existing load balancing rules so that they continue to function.

¹ http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=10051/

Apêndice A. Time Zones

The following time zone identifiers are accepted by CloudStack. There are several places that have a time zone as a required or optional parameter. These include scheduling recurring snapshots, creating a user, and specifying the usage time zone in the Configuration table.

Etc/GMT+12	Etc/GMT+11	Pacific/Samoa
Pacific/Honolulu	US/Alaska	America/Los_Angeles
Mexico/BajaNorte	US/Arizona	US/Mountain
America/Chihuahua	America/Chicago	America/Costa_Rica
America/Mexico_City	Canada/Saskatchewan	America/Bogota
America/New_York	America/Caracas	America/Asuncion
America/Cuiaba	America/Halifax	America/La_Paz
America/Santiago	America/St_Johns	America/Araguaina
America/Argentina/ Buenos_Aires	America/Cayenne	America/Godthab
America/Montevideo	Etc/GMT+2	Atlantic/Azores
Atlantic/Cape_Verde	Africa/Casablanca	Etc/UTC
Atlantic/Reykjavik	Europe/London	CET
Europe/Bucharest	Africa/Johannesburg	Asia/Beirut
Africa/Cairo	Asia/Jerusalem	Europe/Minsk
Europe/Moscow	Africa/Nairobi	Asia/Karachi
Asia/Kolkata	Asia/Bangkok	Asia/Shanghai
Asia/Kuala_Lumpur	Australia/Perth	Asia/Taipei
Asia/Tokyo	Asia/Seoul	Australia/Adelaide
Australia/Darwin	Australia/Brisbane	Australia/Canberra
Pacific/Guam	Pacific/Auckland	

Apêndice B. Event Types

VM.CREATE	TEMPLATE.EXTRACT	SG.REVOKE.INGRESS
VM.DESTROY	TEMPLATE.UPLOAD	HOST.RECONNECT
VM.START	TEMPLATE.CLEANUP	MAINT.CANCEL
VM.STOP	VOLUME.CREATE	MAINT.CANCEL.PS
VM.REBOOT	VOLUME.DELETE	MAINT.PREPARE
VM.UPGRADE	VOLUME.ATTACH	MAINT.PREPARE.PS
VM.RESETPASSWORD	VOLUME.DETACH	VPN.REMOTE.ACCESS.CREATE
ROUTER.CREATE	VOLUME.UPLOAD	VPN.USER.ADD
ROUTER.DESTROY	SERVICEOFFERING.CREATE	VPN.USER.REMOVE
ROUTER.START	SERVICEOFFERING.UPDATE	NETWORK.RESTART
ROUTER.STOP	SERVICEOFFERING.DELETE	UPLOAD.CUSTOM.CERTIFICATE
ROUTER.REBOOT	DOMAIN.CREATE	UPLOAD.CUSTOM.CERTIFICATE
ROUTER.HA	DOMAIN.DELETE	STATICNAT.DISABLE
PROXY.CREATE	DOMAIN.UPDATE	SSVM.CREATE
PROXY.DESTROY	SNAPSHOT.CREATE	SSVM.DESTROY
PROXY.START	SNAPSHOT.DELETE	SSVM.START
PROXY.STOP	SNAPSHOTPOLICY.CREATE	SSVM.STOP
PROXY.REBOOT	SNAPSHOTPOLICY.UPDATE	SSVM.REBOOT
PROXY.HA	SNAPSHOTPOLICY.DELETE	SSVM.H
VNC.CONNECT	VNC.DISCONNECT	NET.IPASSIGN
NET.IPRELEASE	NET.RULEADD	NET.RULEDELETE
NET.RULEMODIFY	NETWORK.CREATE	NETWORK.DELETE
LB.ASSIGN.TO.RULE	LB.REMOVE.FROM.RULE	LB.CREATE
LB.DELETE	LB.UPDATE	USER.LOGIN
USER.LOGOUT	USER.CREATE	USER.DELETE
USER.UPDATE	USER.DISABLE	TEMPLATE.CREATE
TEMPLATE.DELETE	TEMPLATE.UPDATE	TEMPLATE.COPY
TEMPLATE.DOWNLOAD.START	TEMPLATE.DOWNLOAD.SUCCESS	TEMPLATE.DOWNLOAD.FAILED
ISO.CREATE	ISO.DELETE	ISO.COPY
ISO.ATTACH	ISO.DETACH	ISO.EXTRACT
ISO.UPLOAD	SERVICE.OFFERING.CREATE	SERVICE.OFFERING.EDIT
SERVICE.OFFERING.DELETE	DISK.OFFERING.CREATE	DISK.OFFERING.EDIT
DISK.OFFERING.DELETE	NETWORK.OFFERING.CREATE	NETWORK.OFFERING.EDIT
NETWORK.OFFERING.DELETE	POD.CREATE	POD.EDIT
POD.DELETE	ZONE.CREATE	ZONE.EDIT
ZONE.DELETE	VLAN.IP.RANGE.CREATE	VLAN.IP.RANGE.DELETE
CONFIGURATION.VALUE.EDIT	SG.AUTH.INGRESS	

Apêndice C. Alerts

The following is the list of alert type numbers. The current alerts can be found by calling listAlerts.

MEMORY = 0

CPU = 1

STORAGE =2

STORAGE_ALLOCATED = 3

PUBLIC_IP = 4

PRIVATE_IP = 5

HOST = 6

USERVM = 7

DOMAIN_ROUTER = 8

CONSOLE_PROXY = 9

ROUTING = 10// lost connection to default route (to the gateway)

STORAGE_MISC = 11 // lost connection to default route (to the gateway)

USAGE_SERVER = 12 // lost connection to default route (to the gateway)

MANAGMENT_NODE = 13 // lost connection to default route (to the gateway)

DOMAIN_ROUTER_MIGRATE = 14

CONSOLE_PROXY_MIGRATE = 15

USERVM_MIGRATE = 16

VLAN = 17

SSVM = 18

USAGE_SERVER_RESULT = 19

Apêndice C. Alerts

```
STORAGE_DELETE = 20;
```

```
UPDATE_RESOURCE_COUNT = 21; //Generated when we fail to update the resource count
```

```
USAGE_SANITY_RESULT = 22;
```

```
DIRECT_ATTACHED_PUBLIC_IP = 23;
```

```
LOCAL_STORAGE = 24;
```

```
RESOURCE_LIMIT_EXCEEDED = 25; //Generated when the resource limit exceeds the limit.  
Currently used for recurring snapshots only
```

Apêndice D. Revision History

Revisão 0-0 **Tue May 29 2012**

Jessica Tomechak

Initial creation of book by publican

