Stratus Cloud Solution

Installation Guide

Stratus Technologies
For an Always-On World
www.stratus.com

Notice

The information contained in this document is subject to change without notice.

UNLESS EXPRESSLY SET FORTH IN A WRITTEN AGREEMENT SIGNED BY AN AUTHORIZED REPRESENTATIVE OF STRATUS TECHNOLOGIES, STRATUS MAKES NO WARRANTY OR REPRESENTATION OF ANY KIND WITH RESPECT TO THE INFORMATION CONTAINED HEREIN, INCLUDING WARRANTY OF MERCHANTABILITY AND FITNESS FOR A PURPOSE.

Stratus Technologies assumes no responsibility or obligation of any kind for any errors contained herein or in connection with the furnishing, performance, or use of this document. Software described in Stratus documents (a) is the property of Stratus Technologies Bermuda, Ltd. or the third party, (b) is furnished only under license, and (c) may be copied or used only as expressly permitted under the terms of the license.

Stratus documentation describes all supported features of the user interfaces and the application programming interfaces (API) developed by Stratus. Any undocumented features of these interfaces are intended solely for use by Stratus personnel and are subject to change without warning.

Copyrights

Stratus, the Stratus logo, everRun, and SplitSite are registered trademarks of Stratus Technologies Bermuda, Ltd. The Stratus Technologies logo, the Stratus 24 x 7 logo, and Automated Uptime are trademarks of Stratus Technologies Bermuda, Ltd.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Intel and the Intel Inside logo are registered trademarks and Xeon is a trademark of Intel Corporation or its subsidiaries in the United States and/or other countries/regions.

Microsoft, Windows, Windows Server, and Hyper-V are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries/regions.

VMware is a registered trademarks of VMware, Inc. in the United States and/or other jurisdictions.

The registered trademark Linux is used pursuant to a sublicense from the Linux Mark Institute, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Google and the Google logo are registered trademarks of Google Inc., used with permission. The Chrome browser is a trademarks of Google Inc., used with permission.

Mozilla and Firefox are registered trademarks of the Mozilla Foundation.

Red Hat is a registered trademarks of Red Hat, Inc. in the United States and other countries.

Dell is a trademark of Dell Inc.

Hewlett-Packard and HP are registered trademarks of Hewlett-Packard Company.

All other trademarks and registered trademarks are the property of their respective holders.

Manual Name: Stratus Cloud Solution Installation Guide

Product Release Number: Stratus Cloud Solution Release 1.5.1.0

Publication Date: Friday, May 08, 2015

Stratus Technologies, Inc.

111 Powdermill Road

Maynard, Massachusetts 01754-3409

Table of Contents

1

Chapter 1: Stratus Cloud Solution Installation Guide

This installation guide does not intend to provide a one-stop solution to all the issues in a production-grade cloud, but serves as guide to be used as a baseline for building the production OpenStack Stratus Cloud Solution.

This guide also provides technically knowledgeable field engineers with an option to set up a production-grade OpenStack cloud with all the features that are required for and supported by the Stratus Cloud Solution. The architecture of this implementation can be extended or simplified to meet your requirements.

Related Topics

"Stratus Cloud Solution Installation Overview" on page 2

"OpenStack Overview" on page 2

"Stratus Cloud Solution Hardware Overview" on page 4

"Stratus Cloud Solution Software Overview" on page 5

"OpenStack Installation and Configuration" on page 8

"OpenStack Verification and Testing" on page 23

"Installing the Stratus Cloud Solution" on page 25

"Installing Stratus Cloud Availability Services" on page 52

Stratus Cloud Solution Installation Overview

The Stratus Cloud Solution enables IT administrators to efficiently implement and manage a multiple availability level cloud, including support for highly available legacy applications. This allows IT administrators to provide an intuitive service catalog and application administration for end users.

OpenStack installation and configuration is not always an easy task. Each OpenStack installation is unique by virtue of the workload it intends to support, the differences in hardware and networking, and the security and compliance requirements; however, in each case, you install Stratus Cloud Workload Services in a CentOS virtual appliance, which you deploy as an instance on a KVM hypervisor in your OpenStack cloud.

Related Topics

"OpenStack Overview" on page 2

OpenStack Overview

OpenStack is a group of interrelated open-source projects designed to provide massively scalable public and private clouds. The following services and projects are used throughout this document.

Service	Project Name	Description
Dashboard	Horizon	Allows you to interact with OpenStack services to launch an instance, assign IP addresses, set access controls, and other parameters.
Compute	Nova	Provisions and manages large networks of virtual machines on demand.
Networking	Neutron	Enables network connectivity as a service among interface devices managed by other OpenStack services; usually Compute. Allows you to create and attach interfaces to networks. Neutron features a plugable architecture that supports many popular networking vendors and technologies.
Object Storage	Swift	Stores and gets files. Does not mount directories like a file server.

Service	Project Name	Description
Block Storage	Cinder	Provides persistent block storage to guest virtual machines.
Identity service	Keystone	Provides authentication and authorization for the OpenStack services. Also provides a service catalog within a particular OpenStack cloud.
Image service	Glance	Provides a registry of virtual machine images; used by Compute to provision instances.
Telemetry service	Ceilometer	Monitors and meters the OpenStack cloud for billing, benchmarking, scalability, and statistics purposes.
Orchestration service	Heat	Orchestrates multiple composite cloud applications by using either the native HOT template format or the AWS CloudFormation template format, through both an OpenStack-native REST API and a CloudFormation-compatible query API.

Stratus Cloud Solution Hardware Overview

This section provides a general overview of hardware requirements and recommendations. Your exact hardware requirements should be calculated by the number of instances and resource needs of the workloads.

Related Topics

"Stratus Cloud Solution Hardware Requirements" on page 4

Stratus Cloud Solution Hardware Requirements

The Stratus Cloud Solution requires an OpenStack environment with a minimum of:

- One OpenStack controller running Horizon, Keystone, Glance, and Heat

  OpenStack projects can be configured to run on separate nodes; however, for simplicity and maintainability, the core projects (Horizon, Keystone, and Glance) are typically deployed on a single node. Although Heat is usually optional, the Stratus Cloud Solution requires it for orchestration services.

- Two Nova compute nodes running the KVM hypervisor

  Compute nodes provide CPU and memory for your instances. When calculating your hardware requirements, consider the expected number of instances and resource needs of your workloads.

- One Neutron network node

  Neutron is typically configured to run on a separate node because it requires high bandwidth, but it may be installed on the OpenStack controller for smaller cloud configurations.

- Optionally, one Cinder storage node

  Block storage (Cinder) is configured on separate node clusters with expandable storage.

  Cinder resiliency must be supplied by your Cinder storage solution, especially if you deploy mission-critical applications.

  A 10Gb storage network is recommended for best performance and reliability.

Stratus Cloud Solution Software Overview

This section provides a general overview of software requirements, versioning, and recommendations.

Related Topics

"Required OpenStack Components" on page 5

"Supported Internet Browsers" on page 5

"Supported Guest Operating Systems" on page 6

Required OpenStack Components

The Stratus Cloud Solution requires a fully installed and functioning installation of OpenStack Icehouse 2014.1.3 (Version 4) on CentOS 6.6.

At a minimum, you must install the following OpenStack components :

- Keystone (identity service)
- Glance (image service)
- Nova (compute service)
- Neutron (networking service)
- Heat (orchestration service)

For additional requirements related to installing Stratus Cloud Availability Services, see "KVM-FT Hardware and Software Requirements" on page 53.

Supported Internet Browsers

The following table lists the supported operating systems and browsers.

Browsers

Browser	Version
Firefox (on Linux only)	27
Internet Explorer (Windows)	11
Firefox (Windows; Technology Preview only)	27
Google Chrome (Windows and Linux; Technology Preview only)	32.0

Supported Guest Operating Systems

The following operating systems have been tested for deployment into a Stratus-managed OpenStack environment.

Windows-based Operating Systems

Operating System	Version
Windows Server 2008	R2
Windows Server 2012	R2

Linux-based Operating Systems

Operating System	Version
CentOS 64-bit	6.5, 6.6, 7.0
Linux Ubuntu	14.04
Fedora	20, 21
SUSE	12.3

Sample Network Design

The following table shows the networks configured in this implementation. This installation further assumes that the networks are configured using VLANs on the bonded *eth1* interface, and that *eth0* is the administrative network.

Purpose	Is routed?	Network Name	Examples
Provides IPMI and SSH connectivity. This network is only available within the data-center.	Yes	Admin net	192.168.81.0/24

Purpose	Is routed?	Network Name	Examples
Routed network to access OpenStack APIs.	Yes	API net	192.168.82.0/24
Routed network that provides connectivity to the floating IPs assigned to the virtual machines.	Yes	External net	192.168.83.0/24
OpenStack internal messageQ communication non-routed network.	No	MGMT net	10.200.10.0/24
OpenStack virtual machine to virtual machine communication non-routed network.	No	Data net	10.10.1.0/24
OpenStack cinder and NFS storage non-routed network.	No	Cinder storage	10.200.11.0/24

Provider Networks

Provider networks allow cloud administrators to create OpenStack Networking networks that map directly to physical networks in the data center. The Provider networks are the backbone of a virtual network(s) which maps virtual network to physical network. In the Stratus, the provider network is the externally-routed network.

OpenStack Installation and Configuration

This section provides information for installations and configurations required for your OpenStack and Stratus Cloud Solution systems.

> **Note**: Stratus labs are currently using CentOS 6.6. Unless otherwise stated, assume CentOS 6.6 as the operating system running on the systems.

Related Topics

"Physical Connectivity" on page 8

"Storage Considerations" on page 9

"Managing OpenStack Quotas" on page 9

"Managing RabbitMQ File Descriptor Limits" on page 13

"Configuring Instances to Start Automatically in OpenStack" on page 14

"Configuring OpenStack for Evacuation and Migration" on page 16

"Disabling Checksum Offload on OpenStack Component Nodes" on page 19

"GRE Network Configuration" on page 21

"MySQL Connections" on page 21

Physical Connectivity

In order to achieve highest level of availability, the physical network connectivity must be redundant. The following is a non-exhaustive list of suggestions recommended by Stratus:

- Hardware configuration should include at least two NICs, with enough ports in the card to support your networking requirements.
- Bond two ports from two different physical cards to protect against NIC port and NIC failure.
- Connect each port of the server bonded network to separate switches to protect against switch failures.

In the following diagram, the server contains two network cards; each card has two ports. The ports are marked as Card1Port1 through Card2Port2. Bond Card1Port1 and Card2Port1, and connect the ports to two separate switches.

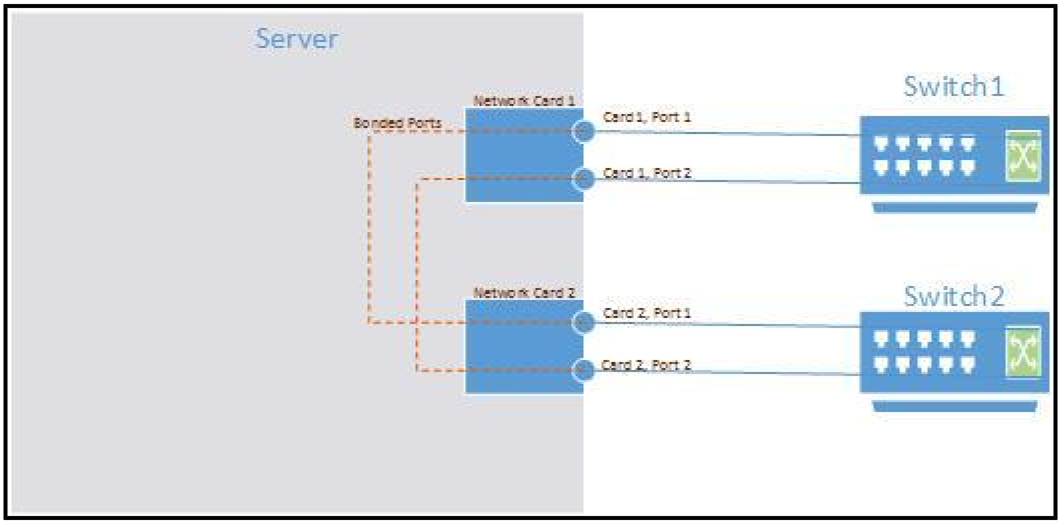

For the switch configuration, configure VLANs on the switch matching a similar configuration to the one in "Sample Network Design" on page 6. Assign a temporary IP address, and verify network connectivity.

Storage Considerations

Stratus recommends the following general considerations when setting up your storage environment.

- Protect the boot volume with either RAID 5 or RAID 6. A single disk boot volume can result in system failure (and failure of all instances) if the disk fails.
- Shared storage should be highly available, as all instances live on the shared storage.
- Cinder block shared storage should be highly available. Unless a high-end storage solution is used in the backend, the disks must at least be RAID 5 or 6 for physical disk failures protection.
- Use a 10Gb storage network for best performance and reliability.

Managing OpenStack Quotas

To prevent OpenStack system resources from being exhausted, modify the OpenStack quotas that you set for each tenant to account for the additional demands of Stratus Cloud Workload Services.

OpenStack enforces quotas for services including:

- Nova (compute) service, to control the number of instances and the number of cores and amount of RAM available to each tenant's applications.
- Neutron (networking) service, to control the number of networks, ports, and subnets.

- Cinder (storage) service, to control the amount of storage space and number of volumes and snapshots per tenant.
- Heat (orchestration) service, to control the maximum number of stacks per tenant and resources per stack.

For each of these services, you must configure reasonable quota settings to account for the applications that you will deploy as well as the additional resources necessary for managing the applications in Workload Services. Some examples follow.

Notes:

1. For best results, set quotas by using OpenStack command-line utilities instead of OpenStack Horizon, which may not properly enforce your settings.
2. There is some overlap between the quotas enforced by the Nova and Neutron services. In the event of a conflict, Neutron quotas always take precedence.
3. If you are uncertain about how much to increase quotas, as a general rule consider increasing quotas by a factor of 10. For example, if a default quota is 10, multiply by a factor of 10 and set the quota to 100.

Increasing Nova Service Quotas

Stratus recommends increasing Nova service quotas including the limits for security groups and security group rules available to your cloud. Use the following command to display the current quota settings for the Nova service, which usually runs on your OpenStack controller:

```
# nova quota-show
+-----------------------------+-------+
| Quota                       | Limit |
+-----------------------------+-------+
| instances                   | 10    |
| cores                       | 20    |
| ram                         | 51200 |
| floating_ips                | 10    |
| fixed_ips                   | -1    |
| metadata_items              | 128   |
| injected_files              | 5     |
| injected_file_content_bytes | 10240 |
| injected_file_path_bytes    | 255   |
| key_pairs                   | 100   |
| security_groups             | 10    |
| security_group_rules        | 20    |
+-----------------------------+-------+
```

For example, the default Nova quotas allow a tenant to create only 10 security groups and 20 security group rules, but even a single tenant could easily exceed those limits, especially if you create deployment packages with multiple instances and multiple internal/external connection rules. You must adjust the quotas according to your needs.

Use the `nova quota-update` command to change quota values per tenant. For example, enter a command similar to the following to change the `security_groups` quota value:

```
# nova quota-update --tenant_id 1cc75c171d764b5584c70be22bf53105 --
security_groups 100
```

If you do not want to manage the individual Nova quotas, you can also disable them by setting the value of each quota to `-1`.

Update the quota values only for tenants that will run applications, including the **Default Tenant**. To determine the `tenant_id` to specify when updating quotas, run the **keystone tenant-list** command, as follows:

```
# keystone tenant-list
+----------------------------------+-----------------------+------
---+
|                id                |          name         |
enabled |
+----------------------------------+-----------------------+------
---+
| 33711efdce85470ca3b8097138207597 | CloudMgmt.Orchestrator |
True  |
| e9b8e27e6ece48249a37f8497df20fe5 |    CloudMgmt.Users     |
True  |
| 1cc75c171d764b5584c70be22bf53105 |     Default Tenant      |
True  |
| 5e9fdbab64204b84a8a6b67994516ab5 |         admin           |
True  |
| 2fc6692d4c1d406cb0fb96d21cf8ef43 |          demo           |
True  |
| bf184899dcfa49b1a31b0f0ae8f96c10 |        service          |
True  |
+----------------------------------+-------------------------+------
---+
```

Increasing Neutron Quotas

Stratus recommends increasing Neutron service quotas including the limit for number of networks available to your cloud. Use the following command to display the current quota settings for the Neutron service, which usually runs on your OpenStack controller:

```
# neutron quota-show
+--------------------+
| Field      | Value |
+------------+-------+
| floatingip | 20    |
| network    | 5     |
| port       | 20    |
| router     | 10    |
| subnet     | 5     |
+------------+-------+
```

Use the `neutron quota-update` command to change quota values per tenant. For example, enter a command similar to the following to change the `network` quota value:

```
# neutron quota-update --tenant_id 1cc75c171d764b5584c70be22bf53105
--network 50
```

If you do not want to manage the individual Neutron quotas, you can also disable them by setting the value of each quota to `-1`.

Update the quota values only for tenants that will run applications, including the **Default Tenant**. To determine the `tenant_id`, see the instructions under **Increasing Nova Service Quotas**.

Increasing the Number of Stacks

Stratus also recommends increasing the number of stacks available to each tenant in your cloud. You can view the current Heat quotas, including the value of `max_stacks_per_tenant`, in the `/etc/heat/heat.conf` file on your OpenStack controller:

```
# Maximum number of stacks any one tenant may have active at
# one time. (integer value)
#max_stacks_per_tenant=100
```

Because Workload Services creates an environment and application stack for each application, the `max_stacks_per_tenant` default value of 100 limits each tenant to 50 applications. To increase the value to 500, which allows the deployment of up to 250 applications per tenant, execute commands similar to the following:

```
# openstack-config --set /etc/heat/heat.conf DEFAULT max_stacks_
per_tenant 500
# openstack-service restart heat
```

These commands increase the `max_stacks_per_tenant` quota (for all tenants) and restart all heat services to apply the change.

Increasing Other Quotas

Ensure that you increase the quotas for other OpenStack services in a similar manner depending on your needs. For more information about managing quotas, see the OpenStack documentation.

Managing RabbitMQ File Descriptor Limits

If you configure OpenStack services to use RabbitMQ for messaging, you must modify the settings on each RabbitMQ node to increase the maximum number of open files or file descriptors.

Many Linux distributions have a soft limit of 1024 files and a hard limit of 4096 files for each user; however, these numbers can be too low for using RabbitMQ in a production environment. Even reaching the soft limit of 1024 open files can raise an alarm in the operating system that prevents the RabbitMQ service from accepting new connections until the alarm clears, which prevents OpenStack services from functioning properly.

To prevent these problems in production environments, the RabbitMQ documentation recommends increasing the maximum number of files for the `rabbitmq` user to at least 65536 files.

To increase the maximum number of files for RabbitMQ:

1. Log on to a RabbitMQ node.
2. Open or create the `/etc/default/rabbitmq-server` file.
3. Add or edit the following line in the file to increase the file descriptor limit to at least 65536 files:

   ```
   ulimit -n 65536
   ```

4. Save the file and close the text editor.
5. Restart the RabbitMQ service by entering the following command:

   ```
   # service rabbitmq-server restart
   ```

6. Confirm that the new limits are in place by entering the following command:

   ```
   # grep -e Limit -e files /proc/$(cat /var/run/rabbitmq/pid)
   /limits

   Limit                Soft Limit    Hard Limit     Units

   Max open files       65536         65536          files
   ```

7. Repeat the preceding steps to increase the limits on any additional RabbitMQ nodes.

Configuring Instances to Start Automatically in OpenStack

If you want to ensure that all of the instances on a Nova compute node resume their state (for example, start automatically) each time the compute node boots or restarts after a power outage, consider modifying the settings described in this topic.

Cautions:

1. Because the Stratus appliance that supports Workload Services runs as an instance on a Nova compute node in your environment, consider enabling instances to resume their state on the compute node that hosts the Stratus appliance; otherwise, Workload Services will be unavailable until you manually restart the Stratus appliance instance.
2. If a power outage occurs, Workload Services automatically resumes the state of any instances that it manages regardless of the Nova settings discussed in this topic; however, Workload Services cannot manage its instances until the Stratus appliance starts.

To configure instances to start automatically on a Nova compute node:

1. Log on to the console of a compute node as `root` user, or be prepared to use `sudo` to run commands as `root`.
2. Open the `/etc/nova/nova.conf` file.
3. Locate the `resume_guests_state_on_host_boot` property and remove the hashtag (#) at the beginning of the line to uncomment it:

   ```
   # Whether to start guests that were running before the host
   # rebooted (boolean value)
   #resume_guests_state_on_host_boot=false
   ```

4. Set the value of the property to `true`:

   ```
   resume_guests_state_on_host_boot=true
   ```

5. Save the file and close the text editor.
6. If you want to ensure that instances shut down instead of suspending upon shutdown, do the following:

   Caution: If you are modifying settings on the Nova compute node that hosts the Stratus appliance, complete the following steps to ensure that the Stratus appliance shuts down (and restarts) properly.

a. Open the `/etc/init.d/libvirt-guests` file.

b. Locate the `ON_SHUTDOWN` property and change the value from `suspend` to `shutdown`, as follows:

```
ON_SHUTDOWN=shutdown
```

c. Execute the following commands to restart services and apply the changes:

```
# service libvirtd restart
# service openstack-nova-compute restart
```

7. If applicable, repeat this procedure on additional compute nodes.

Configuring OpenStack for Evacuation and Migration

Stratus strongly recommends the following OpenStack compute node configuration to ensure the correct operation for evacuations and migrations. If this configuration is not implemented, evacuations and migrations may fail.

To configure OpenStack for proper evacuation and migration:

1. On each compute node where NFS shared storage is used, add the following options to the NFS mount entry in `/etc/fstab`:
   - `auto,lookupcache=none`
   - Example: `/etc/fstab`:
   - # NFS shared storage for instances:
     - `10.200.11.70:/KVMDataStore /var/lib/nova/instances nfs auto,lookupcache=none 0 0`
2. Complete the following steps to ensure that your SSH keys are properly configured. For additional information, refer to the following websites:
   - https://lists.launchpad.net/openstack/msg24036.html
   - https://ask.openstack.org/en/question/10335/ssh-resize/
   - https://macnugget.org/projects/publickeys/

3. Use an existing SSH key for `/root`, or create a new keypair using the following command:
   - `ssh-keygen -t rsa`
4. Enter the file in which to save the key:
   - `/root/.ssh/id_rsa`
5. Enter the passphrase; leave the field blank for no passphrase.
6. Enter the same passphrase again.
   - Your identification (private key) is saved in `/root/.ssh/id_rsa`.
   - Your public key is saved in `/root/.ssh/id_rsa.pub.`
7. The key fingerprint is `55:45:fc:1f:2d:9b:f5:69:6d:03:5d:ef:2b:50:e8:11 root@<server_name>.<domain>.com`
8. The key's randomart image is:

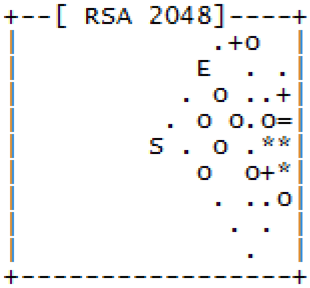

```
+--[ RSA 2048]----+
|            .+o  |
|           E  . .|
|          . o ..+|
|         . o o.o=|
|        S . o .**|
|           o  o+*|
|            . ..o|
|             . . |
|              .  |
+-----------------+
```

9. A key pair is created, both public and private keys:
   - The private root key is located at `/root/.ssh/id_rsa`.
   - The public root key is located at `/root/.ssh/id_ras.pub`.
10. Enable the Nova user for login using the command: `usermod -s /bin/bash nova`
11. Create the folder required by SSH, and move the private key from step 1 into the folder using the following commands:
    - `mkdir -p /var/lib/nova/.ssh`
    - `cp /root/.ssh/id_rsa /var/lib/nova/.ssh`
    - `cat /root/.ssh/id_rsa.pub >> /var/lib/nova/.ssh/authorized_keys`

- Add these to `/var/lib/nova/.ssh/config`
- `Host *`
- `StrictHostKeyChecking no`
- `UserKnownHostsFile=/dev/null`
- `cd /var/lib/nova/.ssh`
- `chown nova *`
- `chgrp nova *`

12. Repeat steps 2 and 3 on each compute node.
13. All nodes share the same key pair; do not generate a new one for the other compute nodes. Instead, copy the key for the compute node on which it was created in step 1. For example:
    - `(copy keys from compute-1 to compute-2) .. scp from compute-2`
    - `scp root@compute-1:/root/.ssh/id_rsa* /root/.ssh`
14. Verify that the key is working properly, using the following commands:
    - `su nova`
    - Example: `ssh nova@compute-1 // you will log in to the node-another without a password`
15. Make sure that all libvirt user IDs and group IDs match across all nodes:
    - For user ID:
      - `id -u qemu`
    - For group ID:
      - `id -g qemu`
      - `id -u nova`
      - `id -g nova`
16. Only on compute nodes that will run the KVM hypervisor, edit `/etc/libvirt/qemu.conf`; uncomment and change these values:

Caution: Do not uncomment or change these values on a compute node that will run the KVM-FT hypervisor.

- `dynamic_ownership=0`
- `user=root`
- `group=root`

17. Reboot the compute node.

Disabling Checksum Offload on OpenStack Component Nodes

Checksum offload (for example, the TCP offload engine (TOE)) improves performance by offloading checksum verification from the system processors to the network interfaces, but it can also lead to data corruption.

To prevent data corruption in your instances, the Stratus Cloud Availability Services installation script automatically disables checksum offload for all network interfaces in your KVM-FT compute nodes. Stratus also recommends that you disable checksum offload on KVM compute nodes that are not running Availability Services as well as other OpenStack components, including your OpenStack controller and any Neutron or Cinder nodes.

Checksum offload is enabled by default in most operating systems, but you can manually disable it as described in the following procedure.

To disable checksum offload for all network interfaces:

1. Log on to the console of a KVM compute node (or other OpenStack component node) as the `root` user, or be prepared to use `sudo` to run commands as `root`.
2. Switch to the `/sbin` directory:

   ```
   # cd /sbin
   ```

3. Search for an existing `/sbin/ifup-local` startup script, which configures network interfaces at startup:

   ```
   # ls ifup-local
   ```

4. If an `/sbin/ifup-local` startup script does not already exist, create a new script as follows:

```
# touch /sbin/ifup-local
# chmod +x /sbin/ifup-local
# chcon --reference /sbin/ifup /sbin/ifup-local
```

The first two commands create the file and mark it as an executable script. The last command sets the SELinux context for the script.

5. Open the `/sbin/ifup-local` startup script in a text editor. Add the following lines to disable checksum offload on all network interfaces (except for the loopback interface) each time the system starts:

```
#!/bin/bash
if [ "$1" != "lo" ]; then
/sbin/ethtool --offload $1 rx off tx off
/sbin/ethtool -K $1 gso off
/sbin/ethtool -K $1 gro off
fi
```

If you are adding the lines to an existing file, the `#! /bin/bash` line is already present. When adding the `if then` statement, ensure that it does not conflict with existing lines in the file.

6. Save and close the `/sbin/ifup-local` startup script.

7. To apply the network startup changes, either restart the system or restart the network.

Caution: Restarting the system or network disrupts any running instances.

For example, to restart the network:

```
# service network restart
```

If you need to avoid restarting for now, you can also manually disable checksum offloading by executing the following `ethtool` command for each network interface :

```
# ethtool -K devname tso off lro off gro off gso off
```

However, the checksum offload features will resume upon restart unless you have updated the `ifup-local` startup script.

8. After disabling checksum offload, verify the status of each network interface, as follows:

   ```
   # ethtool --show-offload devname
   ```

   For example, to verify that all checksum offload settings are `off` for eth0:

   ```
   # ethtool --show-offload eth0
   Features for eth0:
   rx-checksumming: off
   tx-checksumming: off
   scatter-gather: off
   tcp-segmentation-offload: off
   udp-fragmentation-offload: off
   generic-segmentation-offload: off
   generic-receive-offload: off
   large-receive-offload: off
   ntuple-filters: off
   receive-hashing: off
   ```

9. Repeat the preceding steps on additional KVM compute nodes and other OpenStack component nodes, as needed.

GRE Network Configuration

If you are using GRE networking, the maximum transmission unit (MTU) on the appliance must be set to 1400. If you are using VLAN networking, this configuration does not apply.

To configure GRE networking:

1. Edit the script located at `/etc/sysconfig/networking-scripts/ifcfg-eth0`.
2. In the script, change `MTU= "1500"` to `MTU="1400"`.
3. Save the changes in the script.

MySQL Connections

To prevent errors in Heat and Horizon, Stratus recommends setting the number of MySQL connections to 300.

To set MySQL connections:

1. Open the file `/etc/my.cnf`.
2. In the `/etc/my.cnf` file in the `[mysqld]` section, add the parameter: `max_connections = 300`. The following is an example of the the `/etc/my.cnf` file containing the correct MySQL connection settings:

   ```
   [mysqld]
   datadir=/var/lib/mysql
   socket=/var/lib/mysql/mysql.sock
   user=mysql
   # Disabling symbolic-links is recommended to prevent assorted security risks
   symbolic-links=0
   default-storage-engine = innodb
   innodb_file_per_table
   collation-server = utf8_general_ci
   init-connect = 'SET NAMES utf8'
   character-set-server = utf8
   bind-address = 0.0.0.0
   max_connections = 300
   ```

3. Save and close the `/etc/my.cnf` file.
4. Restart `mysqld` using the command `service mysqld restart`.

OpenStack Verification and Testing

OpenStack verification and testing allows you to troubleshoot and resolve issues with your OpenStack and Stratus Cloud Solution configurations. You can identify and resolve issues in:

- Image upload
- Instance creation
- Floating IP configurations
- Security groups
- High Availability testing
- Resiliency testing
- Nodes
- Storage
- Debugging

Related Topics

"Debugging OpenStack" on page 24

"OpenStack Verification" on page 23

"High Availability and Resiliency Testing" on page 23

OpenStack Verification

Use the following procedure for basic OpenStack testing:.

1. Log in to horizon: `controller-1.<yourdomain>.stratus.com`.
2. Upload an image.
3. Create an instance.
4. Create a floating IP.
5. Assign a floating IP to the instance.
6. Create a security group to allow ssh and ping.
7. Verify that you can ping and ssh into the instance.

High Availability and Resiliency Testing

Use the following procedure for High Availability and resiliency testing.

1. Power down one of the switches; then run the OpenStack verification tests described in "OpenStack Verification" on page 23.
2. Power on the switch and power down the other switch, then run the OpenStack verification tests described in "OpenStack Verification" on page 23.
3. Reboot all the nodes in the cluster, then run the OpenStack verification tests described in "OpenStack Verification" on page 23.
4. Storage tests: remove a disk from the RAID5 or RAID6 set. If you have RAID 5, then rebuilding the RAID array may be time-intensive.

See also:

"Debugging OpenStack" on page 24

Debugging OpenStack

Use the following procedures to debug your OpenStack installation and configuration.

1. Start with the controller node to verify that all services are up:

   ```
   $ nova-manage service list

   $ neutron service list

   $ cinder host list
   ```

2. Log files on the controller node are under `/var/log directory`:

   Nova logs: `/var/log/nova`

   Neutron logs: `/var/log/neutron`

   http logs: `/var/log/http`

3. For compute nodes, the logs are under `/var/log/nova`.

Installing the Stratus Cloud Solution

To install the Stratus Cloud Solution:

1. Ensure that your OpenStack environment meets the requirements for installation. See the overview information in the "Stratus Cloud Solution Installation Guide" on page 1.
2. Install a CentOS image to use for the Stratus appliance. See "Installing CentOS" on page 26.
3. If the network firewall at your site prevents you from accessing the Internet, configure the Stratus appliance to work around this restriction for the installation process. See "Preparing to Install the Stratus Cloud Solution Without Internet Access" on page 27.
4. Create an installation configuration file to specify properties for the installation process. See "Creating an Installation Configuration File for Stratus Cloud Workload Services" on page 28.
5. Install the Stratus Cloud Workload Services software in the Stratus virtual appliance. See "Installing Stratus Cloud Workload Services" on page 30.
6. Configure settings as needed in the cloud properties file. See "Configuring the Properties File for the Stratus Cloud Workload Services" on page 32.

   Settings to modify include:

   - "Setting Up Your Mail Server" on page 42
   - "Setting the Logging Detail Level for Stratus Cloud Workload Services" on page 43
7. Configure daily backups of the Workload Services management database that was created during the installation. See "Backing Up the Workload Services Management Database" on page 44.
8. Configure the compute node that hosts the Stratus appliance to resume the state of its instances (for example, start instances automatically) each time the compute node boots or restarts after a power outage; otherwise, Workload Services will be unavailable until you manually restart the Stratus appliance instance. See "Configuring Instances to Start Automatically in OpenStack" on page 14.
9. If applicable, install Stratus Cloud Availability Services on two or more KVM compute nodes. See "Installing Stratus Cloud Availability Services" on page 52.

Related Topics

"Connecting to the Stratus Appliance with SSH" on page 45

"Locating Version Information for Stratus Cloud Workload Services" on page 46

"Upgrading Stratus Cloud Workload Services" on page 47

"Uninstalling Stratus Cloud Workload Services" on page 49

"Stratus Cloud Workload Services Installation Script Options" on page 50

Installing CentOS

The following procedure summarizes how to create and install a custom CentOS image that will serve as the basis for the Stratus appliance that runs Stratus Cloud Workload Services. For a more detailed example of installing a CentOS image, see the OpenStack CentOS image guide.

To install CentOS for Stratus Cloud Workload Services:

1. Upload a CentOS version 6.6 ISO image with the following properties to the `/data/isos` directory of your OpenStack controller:
   - At least 20GB of disk space in the /root partition
   - 4 vCPUs
   - 8 GB RAM
2. Create a virtual hard disk for the image by entering:

   ```
   qemu-img create -f qcow2 /tmp/CentOS-6.6.qcow2 20G
   ```
3. Use `virt-manager` or a similar tool to start the CentOS installation and do the following:
   - Use a custom disk
   - Install cloud-init
   - Create an SSH key using `cloud-user` (or the default). If a password is specified, it must be done from the console, as you cannot SSH to the virtual machine until *cloud-init* is installed. You can use `virt-manager` to access the console.
   - Delete the existing partitions
   - Create a `/root` 20GB partition
   - And `/boot` using the remaining disk space, which is normally about 3 GB
4. After you have finished installing and configuring your CentOS image, upload the image to the OpenStack cloud in Horizon.

5. In OpenStack, create a new flavor for the Stratus appliance instance:
   - 4 vCPUs
   - 16384 MB RAM (no swap,no ephemeral)
   - 160 GB Root Disk
6. Launch an instance for the Stratus appliance and do the following:
   - Enter the instance name, for example **StratusApp**
   - Select the custom flavor that you created for the appliance
   - Select your custom CentOS image
   - Import and select the SSH key associated with the image
   - Select an external network
7. After launching the Stratus appliance, create a configuration file for the installation process as described in "Creating an Installation Configuration File for Stratus Cloud Workload Services" on page 28, and then install the cloud software as described in "Installing Stratus Cloud Workload Services" on page 30.

Preparing to Install the Stratus Cloud Solution Without Internet Access

By default, the Stratus Cloud Workload Services and Availability Services installation scripts require an Internet connection to complete the software installation. If the network firewall rules at your site prevent you from accessing the Internet, you can configure your systems as follows to circumvent the Internet requirement.

To circumvent the Internet requirement for installing Workload Services or Availability Services:

1. Log on to the system where you will be installing the software:
   - To install Workload Services, log on to the Stratus appliance as the `cloud-user` and execute the `su` command to become the `root` user, or be prepared to use `sudo` to run commands as `root`.
   - To install Availability Services, log on to the compute node as the `root` user, or be prepared to use `sudo` to run commands as `root`.
2. Remove the `epel.repo` file associated with the `yum` utility, as follows:

```
# rm -f /etc/yum.repos.d/epel.repo
```

3. Create a blank `epel.repo` file, as follows:

   ```
   # touch /etc/yum.repos.d/epel.repo
   ```

   Creating a blank `epel.repo` file prevents the installation script from creating a version of this file that points to external Internet repositories.

4. Register a different EPEL repository that is available locally at your site.

5. For Availability Services, register additional `yum` repositories at your site that include the CentOS packages and other software packages required to install Availability Services. For a list of the required packages, see "KVM-FT Hardware and Software Requirements" on page 53.

6. For Availability Services, repeat the preceding steps for each compute node on which you will install Availability Services.

Creating an Installation Configuration File for Stratus Cloud Workload Services

Before installing Stratus Cloud Workload Services, you must create a configuration file to specify the settings needed by the installation program. Use the following table to the gather the installation settings.

To create the configuration file, open a text editor in the Stratus virtual appliance and insert a sample configuration file. You can copy the example configuration file that appears below the table or display the help for the installation program (`sudo ./install.sh -h`) and copy the sample file from the output. Paste the content into your text editor, replace the sample settings with the settings for your environment, and save the file (for example, save as `install.conf`).

After creating the configuration file, install the cloud software as described in "Installing Stratus Cloud Workload Services" on page 30. Specify the name of your configuration file to the installation program.

Installation Configuration File Settings

Setting	Description	Label
OpenstackURL	Specify the URL for your OpenStack controller. Find the end point for API access in Horizon, as follows: 1. Click **Compute**, click **Access & Security**, and then click the **API Access** tab. 2. Use the **Identity** service endpoint. For example:	Required

Setting	Description	Label
	`http://192.168.100.50:5000/v2.0`	
OpenstackAdminName	Specify the OpenStack admin account.	Required
OpenstackAdminPassword	Specify the OpenStack admin password.	Required
CloudDbUser	Specify a user name for the main account that Workload Services creates within PostgreSQL to manage the database of cloud configuration data. Only lowercase characters are accepted.	Required
CloudDbPassword	Specify a password for the account specified in the CloudDbUser property.	Required

Example Installation Configuration File

```
# Specify the URL for openstack (REQUIRED)

OpenstackURL=http://192.168.84.50:35357/v2.0

# Specify the admin account (REQUIRED)

OpenstackAdminName=admin

# Specify the admin password (REQUIRED)

OpenstackAdminPassword=admin

# Specify the user for database access, only lowercase characters are accepted. (REQUIRED)

CloudDbUser=clouddbuser

# Specify the password for database access (REQUIRED)

CloudDbPassword=root


# If a specify NTP server or group of NTP servers must be used
# specify one entry per line (OPTIONAL)

NTPServer=192.168.87.150
```

Installing Stratus Cloud Workload Services

Install Workload Services after you have installed CentOS ("Installing CentOS" on page 26) and created an installation configuration file ("Creating an Installation Configuration File for Stratus Cloud Workload Services" on page 28.

To install Workload Services:

1. In Horizon, open the console of the Stratus appliance that you created in the "Installing CentOS" on page 26 procedure. Log on as the *cloud-init* user and enter the password you assigned to *SSH creation* during installation. You may need to update the `/etc/udev/rules.d/70-persistent-net.rules` file and remove the `eth0` entry, and rename `eth1` to `eth0`. This sometimes occurs when you deploy.
2. Update these files as follows:
   - `/etc/sysconfig/network` file: specify the `HOSTNAME` of the appliance. Note that `**NOZEROCONF=yes` should already be there.
   - `/etc/hosts` : add an entry for `127.0.0.1` at the end for the hostname you just updated in the network file.
   - If you have not done so already, install the `cloud-init` using the OpenStack instructions in order to later use the SSH key. Cloud-init software allows an SSH key to be injected to the instance when you launch or deploy an instance. Without cloud-init, you cannot log on using an SSH key.
3. Reboot after these updates.
4. Assign this instance a floating IP address so that you can SSH to it using your key.
5. Log on as the `cloud-user` and execute the `su` command to become the `root` user (or be prepared to use `sudo` to run commands as `root`).
6. Install the `yum-utils` package for access to additional `yum` commands:

   ```
   # yum install yum-utils
   ```

7. Verify that there are no unfinished `yum` transactions by entering the following `yum` command:

```
# yum-complete-transaction --cleanup-only
Loaded plugins: fastestmirror
Loading mirror speeds from cached hostfile
* base: mirror.sanctuaryhost.com
* epel: mirror.cogentco.com
* extras: centos.mirror.nac.net
* updates: mirror.wiredtree.com
No unfinished transactions left.
```

If there are unfinished transactions, resolve them before continuing with the installation.

8. In the `/opt` directory of the Stratus appliance, create a directory called `Release`:

   ```
   # mkdir /opt/Release
   ```

9. Download the Workload Services installation script from the **Stratus Cloud Solution Downloads and Support** page at http://www.stratus.com/services-support/downloads to a local management PC.

10. Transfer the installation script to the `/opt/Release` directory. For example, use a secure copy (SCP) utility to copy the file from the local management PC to the Stratus appliance. (If you copy the script to your `/home/cloud-user` account, move the script to the `/opt/Release` directory.)

11. In the Stratus appliance, switch to the `/opt/Release` directory:

    ```
    # cd /opt/Release
    ```

12. Run the following command to make the installation script an executable file, where *script* is the name of the script:

    ```
    # chmod a+x script.sh
    ```

13. Locate the installation configuration file (for example, `install.conf`) that you created in "Creating an Installation Configuration File for Stratus Cloud Workload Services" on page 28 and move this file to the same `/opt/Release` directory.

14. Run the installation script specifying the `install` option and the name of the installation configuration file; for example:

```
# ./clouds-n.n.n.n.n.sh install install.conf
```

The installation script begins the installation process for Stratus Cloud Workload Services.

> **Notes**:
>
> - For more information about the installation script options, see "Stratus Cloud Workload Services Installation Script Options" on page 50.
> - If the installation fails or you stop the installation script before it can finish, uninstall the cloud software ("Uninstalling Stratus Cloud Workload Services" on page 49), correct any installation errors, and then install the software again. If you retry the installation without uninstalling the software, the script exits with the following error: `There are unfinished transactions remaining. You might consider running yum-complete-transaction first to finish them.`

15. When the installation script is finished, optionally verify the version number of Workload Services that you installed by entering a command similar to the following:

    ```
    # rpm -qa | grep stratus
    stratus-clouds-0.1.5.24.3-0.fc14.noarch
    ```

16. After a successful installation, update the `/opt/jetty/resources/CloudMgmtExt.properties` file to configure settings for Workload Services, as described in "Configuring the Properties File for the Stratus Cloud Workload Services" on page 32.

17. You can now go to `https://ThisCentOSFloatingIP`, and begin using Workload Services.

Related Topics

"Uninstalling Stratus Cloud Workload Services" on page 49

Configuring the Properties File for the Stratus Cloud Workload Services

After installing the Stratus Cloud Workload Services, edit the cloud properties file (`/opt/jetty/resources/CloudMgmtExt.properties`) to configure settings needed for your environment. In most cases, the properties are set automatically during the Workload Services installation, but, for example, you must set up a mail server and optionally set the logging detail level in this file, as described in the following topics:

- "Setting Up Your Mail Server" on page 42
- "Setting the Logging Detail Level for Stratus Cloud Workload Services" on page 43

The table summarizes the settings available in the `CloudMgmtExt.properties` file and indicates if you should modify them in the **Modify** column. An example properties file appears below the table.

> **Note**: If you modify the `CloudMgmtExt.properties` file, you must restart the Jetty service to apply the changes. Restart Jetty by entering the command `service jetty restart`.

Cloud Properties File Settings

Setting	Description	Modify
AllowNonFtOnFt	Specifies if the orchestrator can place standard KVM instances on KVM-FT hypervisors. The default setting of `false` ensures that only KVM-FT instances are placed on KVM-FT hypervisors to reserve the resources and overhead of these hypervisors for **Mission Critical** applications. A setting of `true` allows the orchestrator to place standard KVM instances on KVM-FT hypervisors if hypervisors with a lower **Availability Level** are unavailable or have exceeded their instance load. **NOTE:** The orchestrator cannot evacuate a standard KVM instance if the instance is running on a KVM-FT hypervisor.	No
BootstrapUsername	Specifies the OpenStack admin account.	No

Setting	Description	Modify
BootstrapPassword	Specifies the OpenStack admin password.	No
BootstrapTenantName	Specifies the OpenStack admin tenant name.	No
buggrabber	Specifies the location of the Stratus `bug-grabber` utility that collects log files and other information if needed for your service representative to troubleshoot your cloud configuration.	No
cloud.mgmt.mail.server.host	Specifies the system mail server host to use for email notifications.	Yes
cloud.mgmt.mail.server.port	Specifies the system mail server port, which is 25 by default, and generally 465 for SSL and 587 for TLS.	Yes
cloud.mgmt.mail.sender.address	Specifies the email address that appears in the FROM header when mail is sent.	Yes
cloud.mgmt.mail.secure.password.authentication.require	Specifies if secure password authentication (SPA) is required by email server. **NOTE:** If set to `true`, you must specify the next three properties (encryption type, username, and password).	Yes
cloud.mgmt.mail.encryption.type	Specifies the encryption type of the email server connection (DEFAULT, SSL or TLS). **NOTE:** Must also set cloud.mgmt.mail.secure.password.authentication.required to `true`.	Yes

Setting	Description	Modify
cloud.mgmt.mail.account.username	Specifies the mail account username. **NOTE:** Must also set cloud.mgmt.mail.secure.password.authentication.required to `true`.	Yes
cloud.mgmt.mail.account.password	Specifies the mail account password. **NOTE:** Must also set cloud.mgmt.mail.secure.password.authentication.required to `true`.	Yes
cloud.mgmt.mail.password.reset.url	Specifies the callback URL pattern that is sent by the forgot password email. If necessary, enter the hostname or IP address.	Yes
cloud.mgmt.mail.password.initialize.url	Specifies the callback URL pattern that is sent by the initialize password email. If necessary, enter the hostname or IP address.	Yes
cloud.mgmt.help.system.host	Specifies the URL of the online help system that appears when you click **Help** in Workload Services.	Yes

cloud.mgmt.api.logging.level
cloud.mgmt.orchestrator.logging.level
cloud.mgmt.presentation.api.logging.level
	Specify the logging level to use for Stratus logs in the `/opt/jetty/logs` directory.Defaults are as follows:
	cloud.mgmt.api.logging.level=info
	cloud.mgmt.orchestrator.logging.level=info
	cloud.mgmt.presentation.api.logging.level=info
		Yes

Setting	Description	Modify
	Possible values include: - `fatal`: Shows messages at a FATAL level only - `error`: Shows messages classified as ERROR and FATAL - `warning`: Shows messages classified as WARNING, ERROR, and FATAL - `info`: Shows messages classified as INFO, WARNING, ERROR, and FATAL - `debug`: Shows messages classified as DEBUG, INFO, WARNING, ERROR, and FATAL - `trace` : Shows messages classified as TRACE,DEBUG, INFO, WARNING, ERROR, and FATAL	
databaseUrl	Specifies the URL of the PostgreSQL database.	No
databaseUser	Specifies the PostgreSQL database username.	No
databasePassword	Specifies the PostgreSQL database password.	No
DataCollectionCommodityIntervalCode DataCollectionMissionCriticalIntervalCode DataCol-	Specifies the data collection interval in seconds for each availability type. Valid values are 10, 30, 60, 300, 600, 900, 1800, 2700,	No

Setting	Description	Modify

lectionBusinessCritiCalIntervalCode
DataCollectionDefaultIntervalCode
	and 3600. If an invalid value is set, the software uses the default value. Defaults are as follows:
	DataCollectionCommodityIntervalCode=900
	DataCollectionMissionCriticalIntervalCode=60
	DataCollectionBusinessCritiCalIntervalCode=300
	DataCollectionDefaultIntervalCode=900	

heat.template.function.enable	Used only for internal testing or demonstration. Must always be set to the default of `false`.	No
HypervisorWorkloadUnlimited	Specifies if the cloud allows infinite oversubscription of resources. The default setting of `false` prevents oversubscription of resources. A setting of `true` allows you to continue deploying applications even if cloud resources are exhausted.	
KeystoneEndpoint	Specifies the URL for your OpenStack controller.	No
KvmFtDiskEnableTimeout	Specifies how long the KVM-FT orchestrator will wait (in minutes) before trying to automatically repair an `OFFLINE` or `FAILED` disk. Default is 10 minutes.	Yes
KvmFtQuorum1Ip KvmFtQuorum2Ip	Specifies the IP addresses of the two nodes that run the fault-tolerant Quorum server service (QSS) for KVM-FT applications.	Yes

Setting	Description	Modify
	You manually enter these IP addresses after installing the quorum servers, as described in "Installing Quorum Servers for Stratus Cloud Availability Services" on page 63. You cannot change the quorum server IP addresses later without redeploying your KVM-FT applications.	
MaxAdvisoryDisplay	Specifies the maximum number of records that Workload Services will display within a category of advisories. When trimming to this maximum number, the system displays the most recent advisories. Default value is 500.	Yes
Orchestrator	Used only for internal testing or demonstration. Must remain commented out. Setting to `local` enables a mock orchestrator that simulates the orchestration function without connecting to OpenStack.	No
OrchestratorAPIEndpoint	Specifies the URL of the Heat orchestrator.	No
OSTokenMinutesToRenewBeforeExpires	Specifies renewal interval for OS token in minutes. Default is 15 minutes.	No
SearchHost	Specifies the elastic search path. Default is 127.0.0.1.	No
SearchHttpPort	Specifies the elastic search HTTP port. Default is 9200.	No
SearchPort	Specifies the elastic search port. Default is 9300.	No

Setting	Description	Modify
StackCreateTimeoutMinutes	Specifies how long Workload Services will wait (in minutes) before reporting that application deployment has failed if no prior success or failure response has been received by the system. Default is 60 minutes, as some stack resources take a long time to create (for example, instances with large disks).	Yes

Example Cloud Properties File

```
HypervisorWorkloadUnlimited=false
#Orchestrator=local
KeystoneEndpoint = http://192.168.105.50:5000/v2.0
BootstrapUsername = admin
BootstrapPassword = admin
BootstrapTenantName = admin
# Cloud Management Database Maintenance
StackCreateTimeoutMinutes=60
MaxAdvisoryDisplay=500
SearchHost=127.0.0.1
SearchPort=9300
SearchHttpPort=9200
#minutes to renew before token expires
OSTokenMinutesToRenewBeforeExpires=15
#DATACOLLECTION INTERVAL CODE IN SECONDS
#VALID VALUES ARE (IF INVALID VALUE SET, it WILL DEFAULT to 900)
#10,30,60,300,600,900,1800,2700,3600
DataCollectionCommodityIntervalCode=900
```

```
DataCollectionMissionCriticalIntervalCode=60
DataCollectionBusinessCritiCalIntervalCode=300
DataCollectionDefaultIntervalCode=900
## system mail properties ##
# The system mail server host
cloud.mgmt.mail.server.host=<change_on_setup>
# The system mail server port (25 by default, generally 465 for SSL and 587 for TLS)
cloud.mgmt.mail.server.port=25
# The email address that receiver can see in the header FROM
cloud.mgmt.mail.sender.address=<change_on_setup>
# Set to true if secure password authentication (SPA) is required by email server
cloud.mgmt.mail.secure.password.authentication.required=true
# The encryption type of the email server connection (DEFAULT, SSL or TLS)
cloud.mgmt.mail.encryption.type=DEFAULT
# The mail account username, required if the cloud.mgmt.mail.secure.password.authentication.required
set to true
cloud.mgmt.mail.account.username=<change_on_setup>
# The mail account password, required if the cloud.mgmt.mail.secure.password.authentication.required
set to true
cloud.mgmt.mail.account.password=<change_on_setup>
## password reset mail ##
# The callback URL pattern that sent via the forgot password email. Change the host/ip if necessary
cloud.mgmt.mail.password.reset.url=https://<change_on_setup>/cloud/#-
login:passwordreset?useridentifier={0}&passwordresetcode={1}
## password initialize mail ##
# The callback URL pattern that sent via the initialize password email. Change the host/ip if necessary
```

```
cloud.mgmt.mail.password.initialize.url=https://<change_on_setup>/cloud/#-login:passwordreset?useridentifier={0}&passwordresetcode={1}

# help system host

cloud.mgmt.help.system.host=http://clouddoc.stratus.com/1.5.1.0

#buggrabber=c:/cloud/buggraber.bat

buggrabber=sudo /opt/stratus/scripts/buggrabber.sh

# Orchestration API (currently needed for LAMO trap POSTs)

OrchestratorAPIEndpoint=http://134.111.39.80:8084/orchestrator

# KVM-FT Quorum IPs

KvmFtQuorum1Ip=192.168.81.50

KvmFtQuorum2Ip=192.168.81.90

# KVM-FT disk repair threshold

KvmFtDiskEnableTimeout=10

#### LOGGING LEVEL (OPTIONAL) Default : debug if not found

#POSSIBLE VALUES : DEBUG/INFO/WARN/ERROR/FATAL

#fatal: shows messages at a FATAL level only

#error: Shows messages classified as ERROR and FATAL

#warning: Shows messages classified as WARNING, ERROR, and FATAL

#info: Shows messages classified as INFO, WARNING, ERROR, and FATAL

#debug: Shows messages classified as DEBUG, INFO, WARNING, ERROR, and FATAL

#trace : Shows messages classified as TRACE,DEBUG, INFO, WARNING, ERROR, and FATAL

# APPLIES FOR BOTH API AND ORCHESTRATOR

cloud.mgmt.api.logging.level=debug

cloud.mgmt.orchestrator.logging.level=debug

#Postgresql Database entries

databaseUrl = jdbc:postgresql://localhost:5432/

databaseUser = postgres
```

databasePassword = root

Setting Up Your Mail Server

After installing the Stratus Cloud Workload Services, you must configure the cloud software to connect to the mail server in your environment. In the directory /opt/jetty/resources, edit your CloudMgmtExt.properties file as follows.

> **Note**: Contact your IT department for your specific mail server parameters, which are shown in red in the following example.

To set up your email server:

1. Log on to the Stratus appliance as the cloud-user and execute the su command to become the root user (or be prepared to use sudo to run commands as root).
2. Open the /opt/jetty/resources/CloudMgmtExt.properties file with a text editor and modify the following settings:

```
## system mail properties ##

# The system mail server host

cloud.mgmt.mail.server.host=smtpmail.your_domain.com

# The system mail server port (25 by default, generally 465
for SSL and 587 for TLS)

cloud.mgmt.mail.server.port=25

# The email address that receiver can see in the header FROM

cloud.mgmt.mail.sender.address=user_name@your_domain.com

# Set to true if secure password authentication (SPA) is
required by email server

cloud.mgmt.mail.secure.password.authentication.required=false

# The encryption type of the email server connection (DEFAULT,
SSL or TLS)

cloud.mgmt.mail.encryption.type=DEFAULT
```

```
# The mail account username, required if the cloud.mgmt.mail.secure.password.authentication.required set to true

cloud.mgmt.mail.account.username=<change_on_setup>

# The mail account password, required if the cloud.mgmt.mail.secure.password.authentication.required set to true

cloud.mgmt.mail.account.password=<change_on_setup>

## password reset mail ##

# The callback URL pattern that sent via the forgot password email. Change the host/ip if necessary

cloud.mgmt.mail.password.reset.url=https://your_cloud_server.your_domain.com/cloud/#login:passwordreset?useridentifier={0}&passwordresetcode={1}

## password initialize mail ##

# The callback URL pattern that sent via the initialize password email. Change the host/ip if necessary

cloud.mgmt.mail.password.initialize.url=https://your_cloud_server.your_domain.com/cloud/#login:password?username={0}
```

3. Save and close the `CloudMgmtExt.properties` file.
4. Restart Jetty using the command `service jetty restart`.

Setting the Logging Detail Level for Stratus Cloud Workload Services

After installing Stratus Cloud Workload Services, optionally configure the logging detail level and log file behavior in the Stratus appliance.

To set the logging detail level and behavior for Workload Services:

1. Log on to the Stratus appliance as the `cloud-user` and execute the `su` command to become the `root` user (or be prepared to use `sudo` to run commands as `root`).
2. Open the `/opt/jetty/resources/CloudMgmtExt.properties` file in a text editor and modify the API and Orchestrator logging levels as shown here in red. By default, the logging

detail level is set to `debug`, but you can specify other levels as described in the file:

```
#### LOGGING LEVEL (OPTIONAL) Default : debug if not found

#POSSIBLE VALUES : DEBUG/INFO/WARN/ERROR/FATAL

#fatal: shows messages at a FATAL level only

#error: Shows messages classified as ERROR and FATAL

#warning: Shows messages classified as WARNING, ERROR, and
FATAL

#info: Shows messages classified as INFO, WARNING, ERROR, and
FATAL

#debug: Shows messages classified as DEBUG, INFO, WARNING,
ERROR, and FATAL

#trace : Shows messages classified as TRACE,DEBUG, INFO,
WARNING, ERROR, and FATAL

# APPLIES FOR BOTH API AND ORCHESTRATOR

cloud.mgmt.api.logging.level=info

cloud.mgmt.orchestrator.logging.level=info

cloud.mgmt.presentation.api.logging.level=info
```

3. Save and close the `CloudMgmtExt.properties` file.
4. Open the `/opt/jetty/resources/log4j4OrchExt.xml` file in a text editor and modify the values shown here in red to set the number of log files to back up and the maximum log file size before the system rolls over to a new log file:

```
# APPLIES FOR BOTH API AND ORCHESTRATOR

<param value="20" name="MaxBackupIndex" />

<param value="10MB" name="MaxFileSize" />
```

5. Save and close the `log4j4OrchExt.xml` file.
6. Restart Jetty by entering the command `service jetty restart`.

Backing Up the Workload Services Management Database

Note: For the Stratus Cloud Solution version 1.x, this functionality is for Technology Preview only.

PostgreSQL provides the back-end database for operations in Stratus Cloud Workload Services. When you install Workload Services, the installation script automatically starts the PostgreSQL service in the Stratus virtual appliance and creates a management database. Workload Services populates this database with information about the configuration of your OpenStack environment as well as information about the hypervisors, applications, and instances that you manage with Workload Services.

Caution: Because Workload Services depends on the PostgreSQL database and Stratus appliance for all operations, both must remain running at all times.

In most cases, you do not need to manually manage the database or the appliance, and you must not interfere with their operation; however, it is important to set up daily backups of the PostgreSQL database files to another system to ensure that you can recover your Workload Services configuration in the event of a failure.

Workload Services automatically creates incremental backups of the PostgreSQL database on a daily basis and full backups of the database on a weekly basis. These backups are stored in the `/opt/stratus/backups/database/stratuscloud` directory in the file system of the Stratus appliance. To preserve the database files for recovery, configure your backup server to back up the contents of this directory on a daily basis.

If you ever need to resolve problems with the Stratus appliance or use your database backups to recover the PostgreSQL database, contact your service representative for assistance.

Connecting to the Stratus Appliance with SSH

Open the console of the Stratus appliance if you need to access the command line of the guest operating system for monitoring or troubleshooting purposes.

You can directly open the console of the Stratus appliance in OpenStack Horizon, but if you prefer to connect with a secure shell (SSH) utility, you need to obtain the SSH keys for the instance and specify them to the SSH utility as described in the following procedure.

To access the Stratus appliance, open an SSH connection to the hostname or IP address of the appliance instance. Log on with the default `cloud-user` account, unless you modified this username in your

CentOS image. (If necessary, locate the value for the `default_user` entry in the `/etc/cloud/cloud.conf` file of the image.)

> **Caution**: The Stratus appliance must be running at all times. When accessing the console of the Stratus appliance, be careful not to shut it down or modify its default operating system settings. If the appliance is not working properly, contact your service representative for assistance.

To connect to the console of the Stratus appliance with an SSH utility:

1. Locate the SSH keys that you created when you installed the CentOS operating system.
2. If necessary, transfer the SSH keys to the system where your SSH utility is located. To open the SSH connection, you need only the private key file (not the `.pub` file).
3. Open an SSH connection and log on to the Stratus appliance (as the `cloud-user`):
   - From a Linux-based system, execute the `ssh` command. For example:

     ```
     sudo ssh -i /opt/stratus/keys/StratusCloud.key cloud-user-
     @10.10.10.nn
     ```

   - From a Windows-based system, connect with an SSH utility such as Putty. See the documentation for your SSH utility for information about specifying the SSH key for the connection. For Putty itself, you must use the pre-generated private key file to create a `.ppk` file that is compatible with Putty.

Locating Version Information for Stratus Cloud Workload Services

Locate the version number of Stratus Cloud Workload Services to verify a successful installation or upgrade of the Workload Services software, or to provide this information to your service representative.

To locate the version of Workload Services in the user interface:

1. In the upper right-hand corner of Workload Services, click ⓘ in the **Configuration Panel**.
2. Click **About Stratus Cloud**
3. In the window that appears, note the **Version** number

To locate the version of Workload Services in the console of the Stratus appliance:

1. Log on to the console of the Stratus as the `cloud-user`.
2. Display the version number of Workload Services by entering a command similar to the following:

```
$ rpm -qa | grep stratus
stratus-clouds-0.1.5.24.3-0.fc14.noarch
```

Upgrading Stratus Cloud Workload Services

This topic describes how to upgrade Stratus Cloud Workload Services to a newer version. Use this procedure only to upgrade Workload Services from Version 1.5 to 1.5.x or higher. If you need to upgrade a Version 1.0.x system, contact your service representative for assistance.

Upgrading Workload Services upgrades only the Stratus appliance. If there are associated upgrades for the KVM-FT hypervisors and quorum servers, you must install them separately. See "Upgrading Stratus Cloud Availability Services" on page 71.

To upgrade Workload Services:

1. Log on to Workload Services and verify that your cloud and applications are in a healthy state. Resolve any outstanding advisories before upgrading the software.
2. Log on to the Stratus appliance as the `cloud-user` and execute the `su` command to become the `root` user (or be prepared to use `sudo` to run commands as `root`).
3. Optionally, verify the current version number of Workload Services by entering a command similar to the following:

```
# rpm -qa | grep stratus
stratus-clouds-0.1.5.24.3-0.fc14.noarch
```

4. If you have not already done so, install the `yum-utils` package for access to additional `yum` commands:

```
# yum install yum-utils
```

5. Verify that there are no unfinished `yum` transactions by entering the following `yum` command:

```
# yum-complete-transaction --cleanup-only
Loaded plugins: fastestmirror
Loading mirror speeds from cached hostfile
```

```
* base: mirror.sanctuaryhost.com
* epel: mirror.cogentco.com
* extras: centos.mirror.nac.net
* updates: mirror.wiredtree.com
No unfinished transactions left.
```

If there are unfinished transactions, resolve them before continuing with the upgrade.

6. If you have not already done so, in the `/opt` directory, create a directory called `Release`:

   ```
   # mkdir /opt/Release
   ```

7. Download the new Workload Services installation script from the **Stratus Cloud Solution Downloads and Support** page at http://www.stratus.com/services-support/downloads to a local management PC.

8. Transfer the installation script to the `/opt/Release` directory. For example, use a secure copy (SCP) utility to copy the file from the local management PC to the Stratus appliance. (If you copy the script to your `/home/cloud-user` account, move the script to the `/opt/Release` directory.)

9. In the Stratus appliance, switch to the `/opt/Release` directory:

   ```
   # cd /opt/Release
   ```

10. Run the following command to make the installation script an executable file, where *script* is the name of the script:

    ```
    # chmod a+x script.sh
    ```

11. Ensure that the installation configuration file for the appliance (for example, `install.conf`) is in the same `/opt/Release` directory. Verify that the settings in the configuration file are up to date. For information, see "Creating an Installation Configuration File for Stratus Cloud Workload Services" on page 28.

12. Run the install script specifying the `install` option and the name of the installation configuration file; for example:

    ```
    # ./clouds-n.n.n.n.n.sh install install.conf
    ```

    The installation script begins the upgrade process for the Workload Services software.

> **Notes**:
>
> - The install option automatically detects and upgrades the existing cloud software, if present. For more information about the installation script options, see "Stratus Cloud Workload Services Installation Script Options" on page 50.
> - Do not stop an upgrade while it is running; otherwise, it will leave your cloud in an inconsistent state. If an upgrade fails for any reason, contact your service representative for assistance.

13. When the installation script is finished, optionally verify the new version number of Workload Services by entering a command similar to the following:

```
# rpm -qa | grep stratus
stratus-clouds-0.1.5.25.8-0.fc14.noarch
```

14. Go to https://ThisCentOSFloatingIP and verify that Workload Services is functioning properly.

15. If needed, upgrade your KVM-FT hypervisors. See "Upgrading Stratus Cloud Availability Services" on page 71.

Related Topics

"Uninstalling Stratus Cloud Workload Services" on page 49

Uninstalling Stratus Cloud Workload Services

Uninstall Stratus Cloud Workload Services if the initial installation fails and you need to start over.

> **Note**: To uninstall Workload Services, you must use the installation script for the currently installed version of the software and not the installation script from another build or version. If you followed the installation procedure, the current script is in the /opt/Release directory of the Stratus appliance.

To uninstall Workload Services:

1. Log on to the Stratus appliance as the cloud-user and execute the su command to become the root user (or be prepared to use sudo to run commands as root).

2. In the Stratus appliance, switch to the directory where the existing clouds installation script is

located, typically the `/opt/Release` directory:

```
# cd /opt/Release
```

3. Run the installation script and specify the `uninstall` option:

```
# ./clouds-n.n.n.n.n.sh uninstall
```

4. If applicable, correct any issues that were reported during the initial installation and reinstall the software as described in "Installing Stratus Cloud Workload Services" on page 30.

Stratus Cloud Workload Services Installation Script Options

This topic describes the usage and options of the Stratus Cloud Workload Services installation script.

For an overview of the Stratus Cloud Solution installation procedure, see "Installing the Stratus Cloud Solution" on page 25.

Usage

```
./clouds-n.n.n.n.n.sh subcommand [options]
```

Description

The `clouds` script installs, upgrades, and uninstalls Stratus Cloud Workload Services.

Subcommands

`install` *config-file*	Installs or upgrades Workload Services according to the options in the specified configuration file. To create an installation configuration file, see "Creating an Installation Configuration File for Stratus Cloud Workload Services" on page 28 To reinstall the software with this subcommand, you must uninstall the software first.
`uninstall`	Uninstalls Workload Services. To uninstall the software, you must use the `clouds` script for the currently installed release and not the installation script from another version. If you followed the installation procedure, the current script is in the `/opt/Release` directory of

	the Stratus appliance.
`-h`	Displays generic shell script help.
`-- help`	Displays cloud-specific script help. (You must insert a space between `--` and `help` in this subcommand.)
`--info`	Prints embedded information about the script.
`--list`	Prints a list of files in the installation archive.
`--check`	Verifies the integrity of the installation archive.

Examples

Install or upgrade the software:

```
./clouds-0.1.5.0.0.sh install install.conf
```

Uninstall the software:

```
./clouds-0.1.5.0.0.sh uninstall
```

Installing Stratus Cloud Availability Services

Installing Stratus Cloud Availability Services supplements the standard KVM hypervisor with Stratus availability extensions that allow you to deploy your mission-critical instances in fault-tolerant KVM (KVM-FT) pair groups. You can install Availability Services before or after installing Stratus Cloud Workload Services.

You may prefer to install Availability Services on your compute nodes when you configure your OpenStack environment to meet the requirements of the Stratus Cloud Solution, as described in other topics of the "Stratus Cloud Solution Installation Guide" on page 1.

To install Stratus Cloud Availability Services:

1. Prepare your OpenStack environment for the additional requirements of the KVM-FT hypervisor. See "KVM-FT Hardware and Software Requirements" on page 53.
2. Connect the Ethernet cables between the compute nodes that will run the KVM-FT hypervisor. See "Connecting Ethernet Cables Between KVM-FT Compute Nodes" on page 55.
3. If the network firewall at your site prevents you from accessing the Internet, set up local `yum` software repositories and prepare each compute node to work around this restriction for the installation process. See "Preparing to Install the Stratus Cloud Solution Without Internet Access" on page 27.
4. Install the Availability Services software on the compute nodes. See "Installing Stratus Cloud Availability Services on KVM Compute Nodes" on page 58.
5. Install the Quorum server service (QSS) on two dedicated servers in your OpenStack environment. See "Installing Quorum Servers for Stratus Cloud Availability Services" on page 63.
6. Set the KVM-FT hypervisor pair group(s) and create a test application. See "Configuring the KVM-FT Hypervisor" on page 65.
7. Before deploying any KVM-FT applications for production use, optionally configure your KVM-FT hypervisors and Ethernet switches to increase performance, as described in "Improving the Performance of Stratus Cloud Availability Services" on page 68.

Related Topics

"Locating Version Information for Stratus Cloud Availability Services" on page 70

"Upgrading Stratus Cloud Availability Services" on page 71

"Upgrading Quorum Servers for Stratus Cloud Availability Services" on page 75

"Uninstalling Stratus Cloud Availability Services" on page 77

KVM-FT Hardware and Software Requirements

In addition to the requirements discussed in "OpenStack Installation and Configuration" on page 8, ensure that your OpenStack environment includes:

- At least two dedicated Nova compute nodes that will run the KVM-FT hypervisor

  You install Stratus Cloud Availability Services on these compute nodes, and then select two nodes to be in each fault-tolerant KVM-FT pair group in Stratus Cloud Workload Services. Plan your KVM-FT pair groups carefully, because you cannot change them later without redeploying your KVM-FT applications.

- At least two dedicated nodes that will run the Quorum server service (QSS)

  You install and configure the quorum service on these nodes. The quorum service provides data integrity assurances and automatic restart capabilities for KVM-FT instances. Plan your quorum servers carefully, because you cannot change them later or change their IP addresses without redeploying your KVM-FT applications.

  The quorum service must be installed on servers that are reachable (over UDP) from the KVM-FT hypervisor. Most often, the quorum servers are installed in a similar manner to the OpenStack control plane servers such that the quorum servers are reachable through the OpenStack management network. The quorum protocol generates a small (<1K) packet every second and should not impact regular management traffic.

  The quorum service can be installed on any general-purpose computer with the following attributes:

  - Preferably a bare-metal system
  - Linux-based operating system
  - Small footprint
  - No other services running
  - Memory 512 MB or higher

Be prepared to make the following physical connections for the compute nodes in each KVM-FT pair group:

- Two 10 gigibit (Gb) Ethernet links, known as A-links:
  - Either two 10Gb Ethernet adapters directly connected between the paired compute nodes (minimum configuration) or two 10Gb Ethernet adapters plugged into redundant 10Gb Ethernet switches that connect the paired compute nodes (recommended configuration)
  - Each Ethernet link located on a separate, single-port Ethernet adapter that can be easily replaced

Be prepared to configure storage for KVM-FT applications with:

- Local storage, for best performance and reliability

  KVM-FT instances cannot run on Network File System (NFS) shares because Availability Services requires higher I/O throughput and reliability than NFS can provide. There is also no need to create network shares that can be moved from instance to instance, because Availability Services automatically synchronizes local storage between paired KVM-FT instances in an FT pair group. If a KVM-FT instance fails, Availability Services automatically switches to the secondary instance and its identical local storage.
- Fault-resilient and high-bandwidth Cinder block storage, if needed

  You can use Cinder block storage for KVM-FT instances, but the cloud administrator is responsible for ensuring that the Cinder node is protected against any single points of failure.Also, you should connect the Cinder node to a 10Gb network for best performance and reliability, especially for high I/O applications.

Be prepared to install the following required software packages and repositories during the installation procedure:

> **Note**: You must have Internet access to install the following packages and complete the Availability Services installation. If the network firewall at your site prevents you from accessing the Internet, you must set up local software repositories that contain the following packages and prepare each compute node as described in "Preparing to Install the Stratus Cloud Solution Without Internet Access" on page 27 before beginning the installation procedure.

- Kernel packages for CentOS 6.6:
    - `kernel-2.6.32-504.8.1`
    - `kernel-headers-2.6.32-504.8.1`
    - `kernel-devel-2.6.32-504.8.1`
    - `kernel-firmware-2.6.32-504.8.1`
- GNU Compiler Collection
    - `gcc` (not version specific)
- RPMforge Repository
    - `rpmforge-release-0.5.3-1.el6.rf.x86_64`

Connecting Ethernet Cables Between KVM-FT Compute Nodes

To establish the management network for KVM-FT operations, connect the A-Links between each pair of compute nodes that will run Stratus Cloud Availability Services. There are a few connection options, depending on the level of fault resiliency that you require.

Direct Connection Configuration

For the most basic, minimum configuration, connect two Ethernet cables from 10Gb Ethernet ports on the first KVM-FT compute node to matching Ethernet ports on the second KVM-FT compute node.

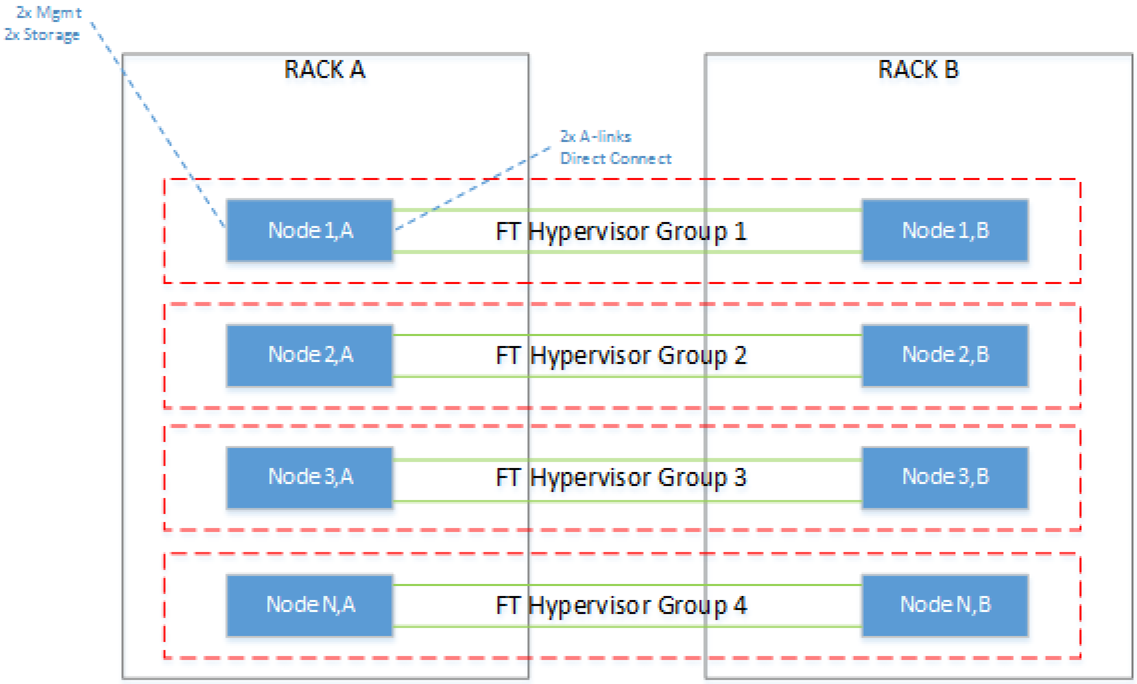

Single Switch Configuration

Alternatively, connect the two KVM-FT compute nodes through a single switch. Connect two Ethernet cables from 10Gb Ethernet ports on the first KVM-FT compute node to a 10Gb Ethernet switch, then connect two additional Ethernet cables from the switch to matching 10 Gb Ethernet ports on the second KVM-FT compute node.

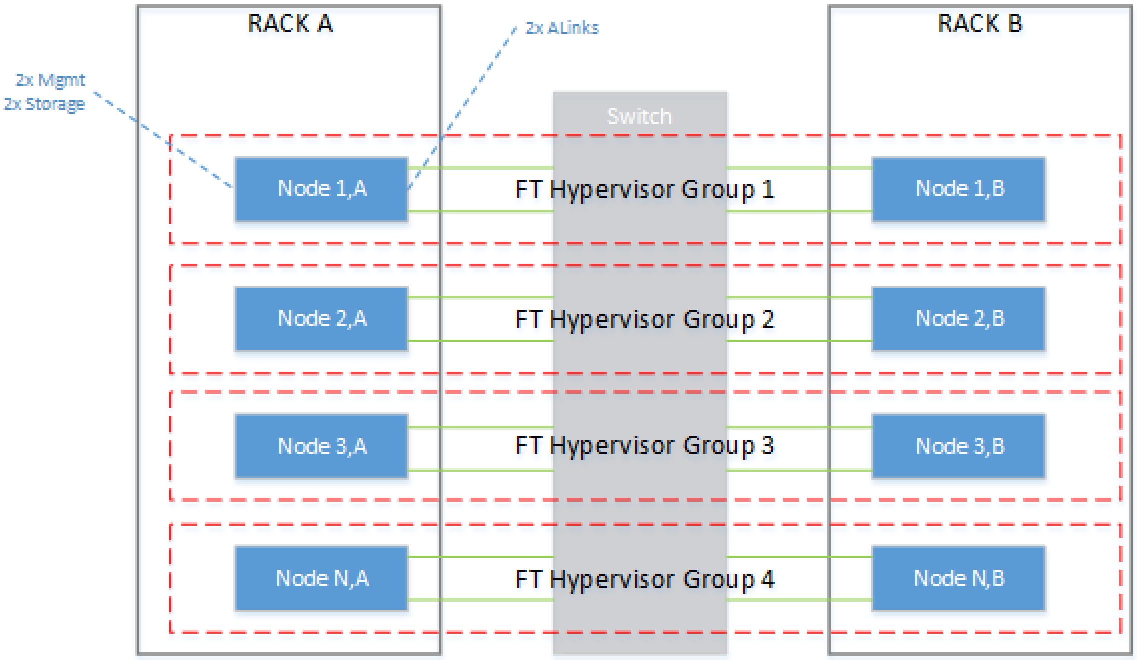

Robust Configuration (Recommended)

For the highest level of fault resiliency, connect the A-Links through a series of redundant 10 Gb Ethernet switches. Even if a cable or switch fails in this configuration, the redundant connections keep your KVM-FT instances running until the problem is corrected.

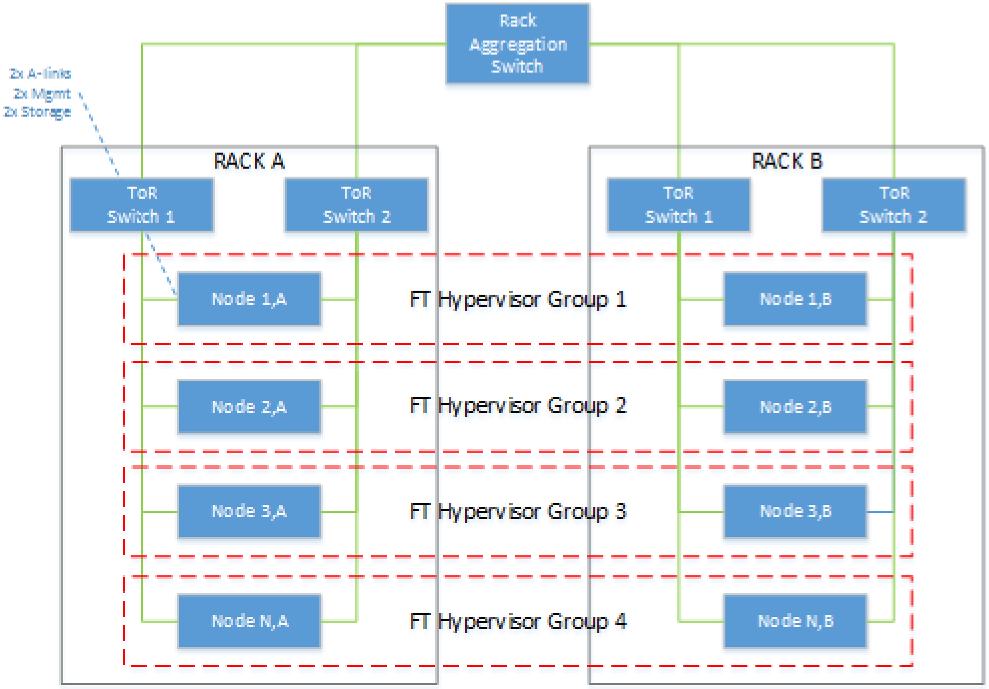

Collecting Ethernet Device Names

Regardless of the configuration you use, make note of the Ethernet device names (for example, `eth1`) for the A-Links on each KVM-FT compute node. You need to specify these device names when you run the Availability Services installation script.

After connecting the cables, install the KVM-FT software as described in "Installing Stratus Cloud Availability Services on KVM Compute Nodes" on page 58.

Installing Stratus Cloud Availability Services on KVM Compute Nodes

Install Stratus Cloud Availability Services after you have prepared your OpenStack environment ("KVM-FT Hardware and Software Requirements" on page 53) and connected the Ethernet cables between the two compute nodes on which you will install the software ("Connecting Ethernet Cables Between KVM-FT Compute Nodes" on page 55).

> **Note**: Because OpenStack Nova sets `yum` priorities that prevent you from installing packages required for Availability Services, you must install perform a clean installation of CentOS 6.6 and upgrade it with the `yum` utility **before** installing OpenStack Nova, as described in this procedure.

To install Availability Services on KVM compute nodes:

1. Perform a clean installation of CentOS 6.6 on a compute node.

2. Log on to the console of the compute node as the `root` user, or be prepared to use `sudo` to run commands as `root`.

3. Run the `yum` utility to update CentOS components including the kernel to the latest available release level, as follows:

   ```
   # yum upgrade
   ```

4. After the upgrade, verify that the CentOS kernel level is `2.6.32-504.8.1` by executing the `uname` command, as follows:

   > **Note**: If yum installs a later kernel version than shown, contact your service representative to verify if the kernel is compatible with Availability Services.

   ```
   # uname -r
   2.6.32-504.8.1.el6.x86_64
   ```

5. Install the additional kernel packages that are required for Availability Services. For example, execute the following command:

   > **Note**: Each kernel package must match the currently installed kernel. As in the following example, if you have installed `kernel-2.6.32-504.8.1`, the correct packages to install are `kernel-`*`package`*`-2.6.32-504.8.1`. It is normal if the `yum` utility reports that one or more of these packages is already installed as a result of the `yum` upgrade.

   ```
   # yum install kernel-headers-2.6.32-504.8.1 \
   kernel-devel-2.6.32-504.8.1 \
   kernel-firmware-2.6.32-504.8.1
   ```

6. Verify that the correct kernel packages are now installed. For example:

   ```
   # rpm -qa | grep kernel
   kernel-devel-2.6.32-504.8.1.el6.x86_64
   ```

```
dracut-kernel-004-356.el6.noarch
kernel-headers-2.6.32-504.8.1.el6.x86_64
libreport-plugin-kerneloops-2.0.9-21.el6.centos.x86_64
kernel-firmware-2.6.32-504.8.1.el6.noarch
abrt-addon-kerneloops-2.0.8-26.el6.centos.x86_64
kernel-devel-2.6.32-504.el6.x86_64
kernel-2.6.32-504.8.1.el6.x86_64
kernel-2.6.32-504.el6.x86_64
```

7. Install the GNU Compiler Collection (`gcc`), as follows:

```
# yum --exclude=*.i686 install gcc
```

8. Add the RPMforge repository to your `yum` configuration, as follows:

```
# yum install http://pkgs.repoforge.org/rpmforge-release/rpmforge-release-0.5.3-1.el6.rf.x86_64.rpm
```

9. Install the `yum-utils` package for access to additional `yum` commands:

```
# yum install yum-utils
```

10. Verify that there are no unfinished `yum` transactions by entering the following `yum` command:

```
# yum-complete-transaction --cleanup-only
Loaded plugins: fastestmirror
Loading mirror speeds from cached hostfile
* base: mirror.sanctuaryhost.com
* epel: mirror.cogentco.com
* extras: centos.mirror.nac.net
* updates: mirror.wiredtree.com
No unfinished transactions left.
```

    If there are unfinished transactions, resolve them before continuing with the installation.

11. To prevent the `ip6tables` firewall service from interfering with KVM-FT hypervisor operations, disable the `ip6tables` service by entering the following commands:

```
# chkconfig ip6tables off
# service ip6tables stop
```

12. In the `/opt` directory, create a directory called `Release`:

    ```
    # mkdir /opt/Release
    ```

13. Download the Availability Services installation script from the **Stratus Cloud Solution Downloads and Support** page at http://www.stratus.com/services-support/downloads to a local management PC.

14. Transfer the installation script to the `/opt/Release` directory. For example, use a secure copy (SCP) utility to copy the file from the local management PC to the compute node.

15. On the compute node, switch to the `/opt/Release` directory:

    ```
    # cd /opt/Release
    ```

16. Run the following command to make the installation script an executable file, where *script* is the name of the script:

    ```
    # chmod a+x script.sh
    ```

17. Run the installation script in the following format, where the *axlinkethdev1* and *axlinkethdev2* arguments represent the Ethernet devices to which the KVM-FT A-link cables are connected and *detectedLink* represents the network (adapter, virtual LAN, or channel bond) that handles data traffic for the instances:

    ```
    # ./kvm-ax-n.n.n.n.n install axlinkethdev1 axlinketdev2 detectedLink
    ```

    The *detectedLink* network is monitored for link up/down state transitions. A link down state indicates that KVM-FT instances may have lost network connectivity for this hypervisor. The KVM-FT software uses the state transitions to determine possible actions.

18. The installation script begins to install the Availability Services software. When prompted, restart the KVM-FT node to apply the changes.

19. After the node restarts, log on and run the `kvmax-healthcheck` script to confirm that the installation was successful. Verify the installed version numbers and examine each section of the

output for any failures. If you need assistance correcting any problems, contact your service representative.

Run the script as follows:

```
# /opt/stratus/scripts/kvmax-healthcheck.sh
- -  - - - - - - - - - - - - - - - - - - - - - - - - - -- -
- - - -
KVM-AX Installer Verification Output
KVM-AX Version : n.n.n.n.n
FTCore Version : n.n.n.n-n.n
- -  - - - - - - - - - - - - - - - - - - - - - - - - - -- -
- - - -
.
.
.
```

20. To prevent the `yum` utility from applying future version upgrades that are incompatible with Availability Services, edit the `/etc/yum.conf` file and ensure that the `exclude` line includes the following entries:

    ```
    exclude=kernel* redhat-release* centos-release* rdo-release*
    ```

    These entries have the following effects:

    - `kernel*` – prevents any kernel package or kernel extension package from being installed
    - `Redhat-release*` – prevents Red Hat from upgrading your operating system version
    - `Centos-release*` – prevents the CentOS release from upgrading your operating system version
    - `Rdo-release*` – controls which version of OpenStack to install

21. Save the changes to the `yum.conf` file.

22. Repeat the previous steps to install the KVM-FT software on the second compute node. If you plan to install the software on more than two compute nodes, you can install the software on all of the

compute nodes now, or continue with installing the quorum servers and configuring your first FT pair group so you can test an application.

23. Install the quorum servers, as described in "Installing Quorum Servers for Stratus Cloud Availability Services" on page 63.

Installing Quorum Servers for Stratus Cloud Availability Services

To maintain the integrity of KVM-FT instances against multiple network failure scenarios, you must install the fault-tolerant Quorum server service (QSS) on two dedicated, Linux-based computers in your OpenStack environment. For information about the system requirements for the quorum servers, see "KVM-FT Hardware and Software Requirements" on page 53.

Cautions:

1. You must install the quorum servers before deploying your first KVM-FT application. Select your quorum servers carefully, because you cannot change them later or change their IP addresses without redeploying your KVM-FT applications.
2. QSS must remain running on both quorum servers to ensure the fault-tolerant operation of your KVM-FT instances. If one node fails, the other node keeps the quorum service running; however, the KVM-FT instances are reported as DEGRADED until the problem is corrected.

To install the quorum servers:

1. Download the quorum service installation script from the Stratus web site.
2. Log on to the console of the first quorum server as the `root` user, or be prepared to use `sudo` to run commands as `root`.
3. In the `/opt` directory, create a directory called `Release`:

   ```
   # mkdir /opt/Release
   ```

4. Use a secure copy (SCP) utility to copy the quorum service installation script to the `/opt/Release` directory.
5. On the quorum server, switch to the `/opt/Release` directory:

   ```
   # cd /opt/Release
   ```

6. Run the following command to make the installation script an executable file, where *script* is the name of the script:

   ```
   # chmod a+x script.sh
   ```

7. Run the installation script specifying the `install` option, as follows:

   ```
   # ./qss-ax-n.n.n.n.n.sh install
   ```

   The installation script installs the quorum service and automatically opens firewall port 4557 (UDP) to allow quorum activity.

8. When the installation script is finished, verify that the quorum service is running by entering the following command:

   ```
   # ps -ef|grep qss
   ```

   If the quorum service is running, it appears in the output as follows:

   ```
   root     25913     1  0 Jan23 ?        00:06:45 /op-
   t/ft/sbin/qss -f
   500      31627 23021  0 18:19 pts/1    00:00:00 grep qss
   ```

9. Record the IP address of the quorum server. For example, execute the `ifconfig -a` command and note the IP address of the `eth0` interface:

   ```
   # ifconfig -a
   eth0      Link encap:Ethernet  HWaddr FA:16:3E:09:32:F8
   inet addr:192.168.101.183  Bcast:192.168.101.255
   Mask:255.255.255.0
   inet6 addr: fe80::f816:3eff:fe09:32f8/64 Scope:Link
   UP BROADCAST RUNNING MULTICAST  MTU:1400  Metric:1
   RX packets:2719992304 errors:0 dropped:0 overruns:0 frame:0
   TX packets:1960011351 errors:0 dropped:0 overruns:0 carrier:0
   collisions:0 txqueuelen:1000
   RX bytes:3721803293741 (3.3 TiB)  TX bytes:1260482520546 (1.1
   TiB)
   ```

```
.
.
.
```

10. Repeat the preceding steps on the second quorum server.

11. After installing the quorum service on both servers, log on to the Stratus appliance as the `cloud-user` and execute the `su` command to become the `root` user (or be prepared to use `sudo` to run commands as `root`).

12. Open the `/opt/jetty/resources/CloudMgmtExt.properties` file in a text editor and locate the `KvmFTQuorum`*n*`IP` entries. Uncomment these entries (remove `#`) and specify the IP addresses of the quorum servers where shown in red:

    ```
    # KVM-FT Quorum IPs
    #KvmFtQuorum1Ip=<change_on_setup>
    #KvmFtQuorum2Ip=<change_on_setup>
    ```

    Enter the IP addresses you collected from the quorum servers. For example:

    ```
    KvmFtQuorum2Ip=192.168.101.183

    KvmFtQuorum1Ip=192.168.101.135
    ```

13. Save and close the `CloudMgmtExt.properties` file.

14. Restart Jetty in the Stratus appliance by entering the command `service jetty restart`.

15. Configure the KVM-FT hypervisor, as described in "Configuring the KVM-FT Hypervisor" on page 65.

Configuring the KVM-FT Hypervisor

After installing the KVM-FT hypervisor ("Installing Stratus Cloud Availability Services on KVM Compute Nodes" on page 58), configure the hypervisors by adding each compute node to a KVM-FT pair group.

> **Caution**: Select your KVM-FT pair groups carefully, because you cannot change them later without redeploying your KVM-FT applications.

> **Note**: You must assign the same **Availability Level** to both hypervisors in an FT pair group; otherwise, applications will fail to deploy. The recommended setting for fault-tolerant operation is **Mission Critical**.

To set a KVM-FT hypervisor pair group:

1. In the main menu of Stratus Cloud Workload Services, click **Hypervisors** to display the **Hypervisors** page.
2. On the **Hypervisors** page, locate the first pair of KVM-FT hypervisors, or compute nodes, that you want to configure as a KVM-FT pair group.
3. Click the first KVM-FT hypervisor that you want to add to the pair group. In the option buttons, click ✏ to display the **Edit Hypervisor** panel.
4. In the **Edit Hypervisor** panel, scroll down to **Service Level Definitions**, and specify the following tags:
   - **Availability Level**: Select **Mission Critical.**
   - **Hypervisor Type**: Select **KVM-FT**.
   - **FT Pair Group**: If this is your first KVM-FT pair group, select **Default Pair Group**. If you are adding another pair group, select a different pair group name. (You can add pair group names as described below under **To add more KVM-FT pair groups**.)

     > **Note**: The **FT Pair Group** field displays only when the selected hypervisor type is KVM-FT.

   - **Hypervisor Instance Storage Type**: Select the storage type for the hypervisor.
   - **Location**: Select one or more deployment locations for the hypervisor.
5. Complete any other fields as needed, and then click **Save**.
6. Repeat steps 3-5 to add the second KVM-FT hypervisor to the same KVM-FT pair group.
7. Ensure that you have installed the two quorum servers required by the KVM-FT hypervisor, as described in "Installing Quorum Servers for Stratus Cloud Availability Services" on page 63.

> **Caution**: The quorum servers must be installed and running before you deploy your first KVM-FT application.

8. Create a test application to verify that your KVM-FT configuration is functioning properly. For an example that summarizes the process of creating an application, from configuring network connections to creating and deploying a service catalog application, see How to Deploy an Application.

9. Before deploying any KVM-FT applications for production use, optionally configure your KVM-FT hypervisors and Ethernet switches to increase performance, as described in "Improving the Performance of Stratus Cloud Availability Services" on page 68.

To add more KVM-FT pair groups (if applicable):

1. Ensure that you have connected the KVM-FT A-links and installed the KVM-FT software on the compute nodes.

2. In the main menu of Stratus Cloud Workload Services, select **Service Level Definitions** to display the **Service Level Definitions** page.

3. On the **Service Level Definitions** page, click **Add Tag** in the upper right corner of the page to display the **Create New Tag** panel.

4. In the **Create New Tag** panel, complete the following:
   - **Name**: Type a name for the new pair group. This name displays on the **Service Level Definitions** page.
   - **Standard Name**: Type a standard name for the new tag. This tag name gets propagated through the back end.
   - **Nest Tag Under**: Select **FT Pair Group**.
   - **Description**: Type descriptive information about the new tag. Information added here displays at the right of the tag name in the listing on the **Service Level Definitions** page.

5. Click **Create**.

6. Repeat steps 1-5 to create additional pair groups, if applicable.

7. To add compute nodes to each new pair group, repeat the procedure above **To set a KVM-FT hypervisor pair group**.

Improving the Performance of Stratus Cloud Availability Services

To improve the performance of Stratus Cloud Availability Services, you can optionally make the following changes to your KVM-FT pair groups. The changes include:

- Enabling *jumbo frames* on each KVM-FT compute node in an FT pair group and the Ethernet switch that connects them.

  Change the *maximum transmission unit* (MTU) from the default of 1500 bytes of payload to the maximum of 9000 bytes. Setting this maximum packet size increases network efficiency by reducing the overall number of packets that each contain networking protocol overhead. (If you want to enable jumbo frames, you must change this setting on each KVM-FT compute and Ethernet switch before the new setting will take effect.)

- Changing a BIOS setting on each KVM-FT compute node to disable the CPU power saving feature and increase CPU performance.

> **Caution**: Because you need to restart each KVM-FT compute node, and potentially restart the Ethernet switch that connects the compute nodes, it is best to make these changes before deploying any KVM-FT applications into production use.

To improve the performance of Availability Services:

1. Log on to the management utility for the Ethernet switch that connects the A-links from your KVM-FT compute nodes. Enable jumbo frames on the switch by setting the MTU size to 9000 bytes. If needed, restart the Ethernet switch to enable the new setting. See the documentation for your Ethernet switch for more information.

2. Log on to the console of the first KVM-FT compute node as the `root` user, or be prepared to use `sudo` to run commands as `root`.

3. Switch to the `/etc/sysconfig/network-scripts` directory:

   ```
   # cd /etc/sysconfig/network-scripts
   ```

4. List the files in the `network-scripts` directory to display the `ifcfg-`*`device`* network startup files (for example, `ifcfg-eth2`). Locate the two `ifcfg-`*`device`* startup files that represent the A-link network interfaces for the compute node.

5. Update the network startup file for each A-link network interface in the compute node to set the MTU value to 9000 bytes. Modify each file as follows:

   a. Open the `/etc/sysconfig/network-scripts/ifcfg-`*`devname`* startup file in a text editor.

   b. If the `MTU` property is already present, change the MTU value to 9000 bytes; otherwise, add a new `MTU=9000` line at the bottom of the file. For example:

   ```
   DEVICE=eth2
   HWADDR=A0:36:9F:2A:BD:3C
   TYPE=Ethernet
   ONBOOT=yes
   NM_CONTROLLED=no
   BOOTPROTO=none
   MTU=9000
   ```

   c. Save and close the network startup file.

   d. Repeat step 4 for the second A-link network interface in the compute node.

6. To apply the network startup changes, restart the compute node, as follows:

   ```
   # shutdown -r now
   ```

   Caution: Restarting the system or network disrupts any running instances.

   Note: Before the operating system starts again, be prepared to press the appropriate key combination to open the system BIOS setup utility.

7. In the system BIOS setup utility for the compute node, locate the **Power Management** properties. If possible, locate a property that controls CPU power and performance. Typically, the default setting throttles CPU performance to save energy. To improve the performance of the KVM-FT compute node, change the value of the property to maximize performance.

   BIOS properties vary by manufacturer and model. See the documentation for your system to locate the correct power-saving property, if present.

8. Save the changes in the BIOS setup utility and restart the compute node.

9. When the operating system starts, log on to the console of the compute node and display the status of each A-link network interface. Verify that the MTU value is set to 9000. For example, to verify the MTU value for `eth2`, execute the following command:

```
# ifconfig eth2
eth2    Link encap:Ethernet  HWaddr A0:36:9F:2A:BD:3C
inet addr:10.200.11.50  Bcast:10.200.11.255
Mask:255.255.255.0
inet6 addr: fe80::225:90ff:fed9:70aa/64 Scope:Link
UP BROADCAST RUNNING MULTICAST  MTU:9000  Metric:1
RX packets:3209 errors:0 dropped:0 overruns:0 frame:0
TX packets:6 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:147614 (144.1 KiB)  TX bytes:468 (468.0 b)
```

10. Repeat the preceding steps on additional KVM-FT compute nodes and Ethernet switches, as needed. Set the same MTU size on both A-links in each KVM-FT compute node of an FT pair group.

Locating Version Information for Stratus Cloud Availability Services

Locate the version number of Stratus Cloud Availability Services to verify a successful installation or upgrade of the Availability Services software, or to provide this information to your service representative.

> **Note**: The Availability Services version number is displayed in Workload Services only if KVM-FT instances are deployed and running.

To locate the version of Availability Services in the Workload Services interface:

1. In the main menu of Workload Services, click **Hypervisors**.
2. On the **Hypervisors** page, select a KVM-FT hypervisor (compute node).
3. Click 🔍 to view the details page for the hypervisor.
4. On the details page:

- Under **KVM-AX Version**, note the current version number of Availability Services.
- If needed, under **FTCore Version**, note the current version number of the FTCore software.

To locate the version of Availability Services in the console of a KVM-FT hypervisor:

1. Log on to the console of a KVM-FT compute node as `root`, or be prepared to use `sudo` to run commands as `root`.
2. Display the contents of the `/etc/nova/kvmax.conf` file. For example:

   ```
   # cat /etc/nova/kvmax.conf
   ```

3. Note the current version numbers in the output. For example:

   ```
   "kvmax-version": "1.5.1.0.16"
   "ftcore-version": "2.0.0.0-21"
   ```

Upgrading Stratus Cloud Availability Services

This topic describes how to upgrade Stratus Cloud Availability Services on your KVM compute nodes. You may need to upgrade Availability Services to install updates specific to the KVM-FT hypervisor or to install new features in conjunction with a new Stratus Cloud Workload Services release.

Upgrade Availability Services on one compute node at a time and in one FT pair group at a time. For example, shut down each KVM-FT instance running on the first compute node and then install the new software release. Start each of the KVM-FT instances on the first node, one instance at a time, and verify that they return to the Running/Paired state. When the first compute node and all of its instances are up and running on the upgraded software, repeat the process on the second node of the FT pair group. In each case, the paired KVM-FT instances on the running compute node keep your applications running during the upgrade.

After upgrading Availability Services, you may also need to upgrade the quorum server service (QSS) on your quorum servers. For more information, see "Upgrading Quorum Servers for Stratus Cloud Availability Services" on page 75.

To upgrade Availability Services on KVM compute nodes:

1. In the main menu of Stratus Cloud Workload Services, click **Hypervisors**.

2. On the **Hypervisors** page, locate the KVM-FT hypervisors (compute nodes). Decide which hypervisor to upgrade first and record its name. Upgrade only one hypervisor at a time, in the same FT pair group.

3. Click the hypervisor that you want to upgrade and click 🔍 to view its details page. On the details page:

   - Under **Deployed Instances**, record the names of the KVM-FT instances that are running on the hypervisor that you selected for upgrade.
   - Under **KVM-AX Version**, note the current version number of Availability Services.

   > **Note**: The version number is displayed only if KVM-FT instances are deployed and running. Version information is also available on each KVM-FT compute node in the `/etc/nova/kvmax.conf` file.)

   When you are finished collecting information about the KVM-FT hypervisor, close the details page.

4. In the main menu, click **Deployed Applications**.

5. On the **Deployed Applications** page, next to **Group by**, click **Availability** and then click **Mission Critical** to expand the category. The **Mission Critical** category typically contains most or all of your KVM-FT instances, but you may need to check other categories as well.

6. Click each application to expand it, and locate the KVM-FT instances that are running on the hypervisor that you want to upgrade. For each application associated with these instances, ensure that one of the instances is in the **Running** state and that the second instance is in the **Paired** state, which indicates that the pair is running normally in fault-tolerant mode.

   > **Caution**: If any of the instances have advisories or are in a state other than Paired/Running, correct any problems before upgrading the hypervisor.

7. If all KVM-FT instances are paired, carefully locate and shut down only the instances on the hypervisor that you want to upgrade. Click the instance, click 🔍, and click ⏻ to stop the instance.

   As long as the application is currently fault-tolerant, and you stop only the instance on the hypervisor you want to upgrade, it does not matter whether that particular instance is in the **Running** or **Paired** state. The other instance keeps the application running.

8. Monitor the instances as they go from the **Stopped** to **Shutoff** state. Ensure that all of the instances are in the **Shutoff** state before you continue. You can monitor the transitions on the **Deployed Applications** page and also on the **Advisories** tab of the **Dashboard** page.

9. When all of the instances on the hypervisor that you want to upgrade are in the **Shutoff** state, log on to the console of the compute node as the `root` user (or be prepared to use `sudo` to run commands as `root`).

10. If you have not already done so, install the `yum-utils` package for access to additional `yum` commands:

```
# yum install yum-utils
```

11. Verify that there are no unfinished `yum` transactions by entering the following `yum` command:

```
# yum-complete-transaction --cleanup-only
Loaded plugins: fastestmirror
Loading mirror speeds from cached hostfile
* base: mirror.sanctuaryhost.com
* epel: mirror.cogentco.com
* extras: centos.mirror.nac.net
* updates: mirror.wiredtree.com
No unfinished transactions left.
```

If there are unfinished transactions, resolve them before continuing with the upgrade.

12. Download the new Availability Services installation script from the **Stratus Cloud Solution Downloads and Support** page at http://www.stratus.com/services-support/downloads to a local management PC.

13. Transfer the installation script to the `/opt/Release` directory. For example, use a secure copy (SCP) utility to copy the file from the local management PC to the compute node.

14. On the compute node, switch to the `/opt/Release` directory:

```
# cd /opt/Release
```

15. Run the following command to make the installation script an executable file, where *script* is the

name of the script:

```
# chmod a+x script.sh
```

16. Run the installation script in the following format, where the *axlinkethdev1* and *axlinkethdev2* arguments represent the Ethernet devices to which the KVM-FT A-link cables are connected and *detectedLink* represents the network (adapter, virtual LAN, or channel bond) that handles data traffic for the instances:

    ```
    # ./kvm-ax-n.n.n.n.n install axlinkethdev1 axlinketdev2 detectedLink
    ```

    The *detectedLink* network is monitored for link up/down state transitions. A link down state indicates that KVM-FT instances may have lost network connectivity for this hypervisor. The KVM-FT software uses the state transitions to determine possible actions.

17. The installation script begins to upgrade the Availability Services software. When prompted, restart the KVM-FT node to apply the changes.

18. After restarting the node, log on and run the `kvmax-healthcheck` script to confirm that the upgrade was successful. Verify the newly installed version numbers and examine each section of the output for any failures. If you need assistance correcting any problems, contact your service representative.

    Run the script as follows:

    ```
    # /opt/stratus/scripts/kvmax-healthcheck.sh
    - -  - - - - - - - - - - - - - - - - - - - - - - - - - - -- -
    - - - -
    KVM-AX Installer Verification Output
    KVM-AX Version : n.n.n.n.n
    FTCore Version : n.n.n.n-n.n
    - -  - - - - - - - - - - - - - - - - - - - - - - - - - - -- -
    - - - -
    .
    ```

.

.

19. In the main menu of Workload Services, click **Deployed Applications**.
20. On the **Deployed Applications** page, next to **Group by**, click **Availability** and then click **Mission Critical** to expand the category.
21. Locate the KVM-FT instances that are on the hypervisor that you just upgraded.
22. Start only one KVM-FT instance and monitor its progress on the **Deployed Applications** page and on the **Dashboard** page as it transitions through various states. It is normal to see critical (red) and warning (orange) states on the **Dashboard** page as startup and sync proceeds.
23. When the **Dashboard** status reaches the normal (green) state and advisories are resolved, go back to the **Deployed Applications** page. For the KVM-FT instance that you started, ensure that the paired instances in the application have returned to the **Running** and **Paired** states.
24. Start the additional KVM-FT instances on the hypervisor that you upgraded one at a time in the same manner.
25. After starting the KVM-FT instances, verify that the new Availability Services version is recognized by Workload Services. On the **Hypervisors** page, click the upgraded hypervisor, click 🔍 to view the details page, and verify the version number under **KVM-AX Version**. (The version is displayed only if KVM-FT instances are deployed and running.)
26. Repeat steps 1-25 for the second hypervisor in the FT pair group.
27. Repeat steps 1-26 for additional FT pair groups.
28. If needed, upgrade the quorum service on your quorum servers.See "Upgrading Quorum Servers for Stratus Cloud Availability Services" on page 75.

Upgrading Quorum Servers for Stratus Cloud Availability Services

This topic describes how to upgrade the Quorum server service (QSS) on the quorum servers associated with Stratus Cloud Availability Services. You may need to upgrade the quorum service as a result of upgrading Availability Services to a new release.

To upgrade the quorum service, first uninstall the QSS package on the quorum server, and then install the new package.

Cautions:

1. Before upgrading the quorum servers, ensure that all of your KVM-FT instances are in the fault-tolerant Running/Paired state.
2. Upgrade one quorum server at a time. Do not upgrade the second quorum server until the first server is running the new software and your KVM-FT instances return to the Running/Paired state.
3. At least one quorum server must be running at all times to maintain the integrity of the KVM-FT instances. It is normal for the KVM-FT instances to become DEGRADED temporarily while a single quorum server is offline for the upgrade.

To upgrade the quorum servers:

1. Log on to Stratus Cloud Workload Services and ensure that all of your KVM-FT instances are in the fault-tolerant Running/Paired state, as follows:
   a. On the **Deployed Applications** page, next to **Group by**, click **Availability** and then click **Mission Critical** to expand the category. The **Mission Critical** category typically contains most or all of your KVM-FT instances, but you may need to check other categories as well.
   b. Click each application to expand it. Ensure that one of the KVM-FT instances is in the **Running** state and that the other instance is in the **Paired** state. If any of the instances have advisories or are in a state other than Running/Paired, correct any problems before upgrading the quorum servers.
2. Log on to the console of the first quorum server as the `root` user, or be prepared to use `sudo` to run commands as `root`.
3. If needed, execute the following command to determine the version of the quorum service that is currently installed.

   ```
   # rpm -qa | grep qss
   lsb-ft-core-cloud-qss-1.0.0.0-112.x86_64
   ```

4. Use a secure copy (SCP) utility to copy the quorum service installation script to the `/opt/Release` directory.
5. On the quorum server, switch to the `/opt/Release` directory:

```
# cd /opt/Release
```

6. Run the following command to make the installation script an executable file, where *script* is the name of the script:

   ```
   # chmod a+x script.sh
   ```

7. Run the installation script specifying the `install` option to upgrade the current installation, as follows:

   ```
   # ./qss-ax-n.n.n.n.n.sh install
   ```

8. Verify that the quorum service is running by entering the following command:

   ```
   # ps -ef|grep qss
   ```

   If the quorum service is running, it appears in the output as follows:

   ```
   root     25913     1  0 Jan23 ?        00:06:45 /opt/ft/sbin/qss -f
   500      31627 23021  0 18:19 pts/1    00:00:00 grep qss
   ```

9. In Workload Services, ensure that all of your KVM-FT instances are in the fault-tolerant Running/Paired state before upgrading the second quorum server.

10. Repeat the preceding steps to upgrade the second quorum server.

Uninstalling Stratus Cloud Availability Services

Uninstall Stratus Cloud Availability Services if you need to upgrade the software as described in "Upgrading Stratus Cloud Availability Services" on page 71, or if the initial installation failed and you need to start over.

> **Note**: To uninstall Availability Services, you must use the installation script for the currently installed version of Availability Services and not the installation script from another build or version. If you followed the installation procedure, the current script is in the `/opt/Release` directory of the KVM-FT compute node.

To uninstall Availability Services:

1. Log on to the console of the compute node as `root`, or be prepared to use `sudo` to run commands as `root`.
2. Delete any KVM-FT applications/instances that you have created.
3. Execute a command similar to the following to uninstall the software:

   ```
   # /opt/Release/kvm-ax-n.n.n.n.n uninstall
   ```

4. If necessary, stop the lamo service. For example, execute:

   ```
   # killall lamo
   ```

5. Repeat these steps on the second KVM-FT compute node.